**IN THE UNITED STATES DISTRICT COURT
FOR THE WESTERN DISTRICT OF TEXAS
WACO DIVISION**

ZENTIAN LTD.,

       Plaintiff,

    v.

APPLE INC.

       Defendant.

C.A. No. 6:22-cv-122

JURY TRIAL DEMANDED

COMPLAINT FOR PATENT INFRINGEMENT

Plaintiff Zentian Ltd. ("Zentian") demands a trial by jury on all issues so triable and, for its complaint against Defendant Apple Inc. ("Apple") alleges as follows:

INTRODUCTION

1.   Apple introduced the "Hey Siri" feature in September 2014.

2.   "Hey Siri" allows users of Apple devices, including iPhone, iPad, Apple Watch, and Mac personal computers, among others, to engage Apple's virtual assistant by using only their voice and without physically handling their device. Apple's virtual assistant does so by recognizing sounds uttered by the device user and converting those acoustic waveforms into information that can be mapped into actions by the operating system, such as searching the internet, playing a song, making to-do lists, setting alarms, creating a calendar invite, or any number of other tasks. This process is known as automatic speech recognition or simply speech recognition.

3.   Apple's implementation of speech recognition that enables users to wake Apple's virtual assistant using the wake phrase "Hey Siri" and/or ask "Siri" questions hands free was invented and patented by Zentian.

4.   Zentian is a technology innovator specializing in hardware designs for speech recognition circuits and methods. Zentian was founded in 2000 by Mark Catchpole using funds from the U.K. government's "Smart Award" program—a program that provides funding for innovative projects with the potential for significant growth.

5.   In the early 2000s, conventional speech recognition systems were ill-equipped to perform a feature like "Hey Siri" on a mobile electronic device in a fast and power-efficient manner, let alone facilitate the back-and-forth conversational interactions that Apple users have with Siri today. In large part, this shortcoming existed because systems in the art focused on improving algorithms and the models used by speech recognition systems to achieve higher word recognition accuracy and other goals. Such implementation focused on incremental improvements in software that ran on power and memory intensive computers or servers. Another contributing factor was the fact that mobile devices like Apple's iPhone or Apple Watch are typically used in noisier environments, such as on public transportation or near a busy road, requiring a more complex model and more intensive computation to obtain adequate results from the speech recognition system.

6.   Zentian, however, focused on speech recognition hardware, including specialized chip designs to enable speech recognition in cost, space, and power constrained environments like mobile electronic devices. Zentian's vision was to study speech recognition systems and to design silicon chips that would accelerate the speech recognition process by off-loading all or part of that processing from software running on a CPU to a bespoke silicon system. This approach differed from that of leading speech recognition companies that were focused on incrementally improving software implementations of speech recognition. Zentian focused on how to design better digital electric circuits (*e.g.*, better hardware).

7.   Critically, speech recognition systems can be divided into several discrete processes. First, an audio front-end generates processed speech parameters (referred to as feature or spectral vectors) that describe the characteristics of an audio signal. Second, a calculating means or circuit compares the similarity between feature vectors and a model of predetermined acoustic states. This stage is often referred to as "distance calculations" because the output is a probability of how similar (or different) a feature vector is from a stored acoustic state. And third, a search process attempts to find the most likely set of words spoken from a given vocabulary based on the distance calculations.

8.   Among the prototypes Zentian developed was a hardware accelerator for the acoustic stage (audio front end and distance calculation) of the speech recognition process that worked alongside a software implementation of the search stage running on a conventional processor. The accelerator demonstrated a two-times speedup of the speech recognition software compared to a conventional speech recognition software-only solution. Moreover, the prototype proved that segregating the speech recognition process onto bespoke hardware components could help solve the shortcomings of conventional systems and enable more advanced features like Apple's "Hey Siri."

9.   In the 2000s, Zentian had a number of conversations about its speech recognition accelerator with leading speech recognition companies, such as Nuance Communications Inc. and IBM, as well as cellular carriers, such as Hutchison 3G UK Limited. But each speech recognition company expressed a desire to focus on software-only solutions and an unwillingness to divert resources to "longer-term" ideas for speech recognition systems. At the 2006 GSM World Congress in Barcelona, a senior executive at Qualcomm, one of the preeminent suppliers of silicon chips for mobile electronic devices, tellingly stated that speech recognition was "not even in my

top ten wish list." Put simply, the industry was not focused on hardware-centric solutions for speech recognition. Zentian was.

10.   Zentian's work resulted in several U.S. patents, including U.S. Patent Numbers 7,587,319; 7,979,277; 10,971,140; 10,062,377; and 10,839,789 (collectively, the "Asserted Patents"). These patents have now been widely cited on patents filed by a host of companies that provide speech recognition technology or end-user products that utilize speech recognition technology, including Samsung, Google, Amazon, Nuance Communications Inc., Microsoft, and Voci Technologies Inc.

11.   Today, speech recognition capabilities are in all kinds of products from phones to cars to smart speakers. Apple's implementation of speech recognition technology in its products "on-device" involves a combination of general-purpose processors combined with hardware accelerators—the very speech recognition circuit and/or system Zentian invented and claimed in the Asserted Patents.

12.   Although Zentian has met with Apple on more than one occasion to discuss its technology, Apple has no license to use Zentian's patented inventions.

13.   Nevertheless, Apple's infringing on-device speech recognition features have, on information and belief, generated substantial sales of Apple products. For example, in the 12 months ending on September 25, 2021, Apple reported over $191 billion in revenue for sales of iPhone, over $35 billion in revenue for sales of Mac computers, over $31 billion in revenue for sales of iPad, and over $38 billion in revenue for sales of wearables, Home and Accessories.[1] On average, end users make multiple requests for Siri per day through "Hey Siri." In turn, Siri and the

---

[1] Apple Inc., *Condensed Consolidated Statements of Operations,* Q4, Fiscal Year 2021, https://www.apple.com/newsroom/pdfs/FY21_Q4_Consolidated_Financial_Statements.pdf.

fluidity of the speech recognition processes that are behind it (facilitated by the hardware architecture on which it runs) have become a material driver of user demand for Apple products.

14.   Apple has infringed and is still infringing, directly and indirectly, one or more claims of the Asserted Patents by making, using, selling, and offering to sell in the United States, including in this District, and importing into the United States, products that implement the Asserted Patents' speech recognition technology. Examples of potentially infringing products include Apple's iPhone, iPad, Watch, AirPods, HomePod, HomePod Mini, Apple TV set top box, and Mac computers (including the Mac Pro, which is manufactured in Austin, Texas), in their various models and iterations (the "Accused Products").[2]

15.   This list of Accused Products is non-limiting and based on information currently available to Zentian. Zentian reserves the right to modify the list of Accused Products as discovery progresses, including as new products are released during the pendency of this case.

THE PARTIES

16.   Zentian is a private corporation incorporated and registered in the United Kingdom, with its registered address at 27 Queens Close, St. Ives, Cambridgeshire, England, PE37 5QD.

17.   Zentian is the owner of all rights, title, and interest in the Asserted Patents.

18.   On information and belief, Defendant Apple Inc. is a California corporation with a principal place of business at 1 Apple Park Way, Cupertino, California. Apple designs, manufactures, and sells throughout the world a wide range of products, including mobile devices and computers that incorporate Zentian's patented technologies.

---

[2] *See, e.g.* Devices that support "Hey Siri" https://support.apple.com/en-us/HT209014. In addition to recognition of the wake phrase "Hey Siri," other potentially infringing features include dictation of text messages and notes with integrated applications and on device speech recognition for Siri user requests.

JURISDICTION AND VENUE

19.   This is an action for patent infringement arising under the patent laws of the United States of America, Title 35 of the United States Code.

20.   This Court has subject matter jurisdiction over the matters asserted in this Complaint under 28 U.S.C. §§ 1331 and 1338(a) and 35 U.S.C. § 271 *et seq*.

21.   This Court's exercise of personal jurisdiction over Apple complies with both the Texas long-arm statute and the Due Process Clause of the Fourteenth Amendment. Apple has purposely directed its activities toward Texas and has availed itself of the privileges of conducting activities in this state.

22.   Apple conducts continuous and systematic business in this District, including by selling infringing products to residents of this District, by soliciting business from residents in this District, and on information and belief, by manufacturing certain Accused Products in this District.

23.   Apple also maintains offices in Austin, Texas, and markets, offers, manufactures, and distributes the Accused Products throughout the United States, including in the state of Texas and in this judicial District.

24.   Apple's Austin offices are located at 12545 Riata Vista Cir, Austin TX 78727.

25.   Venue is proper under 28 U.S.C. § 1391(b) and (c) and 28 U.S.C. § 1400(b).

26.   Apple has committed acts of infringement in this District and has a regular and established place of business here, including offices in Austin, Texas, where it maintains a workforce of more than 7,000 employees.[3]

27.   Apple has maintained a physical presence in this District for years, opening its first offices in Austin in the 1990s. And Apple continues to grow its footprint in this District with a

---

[3] *See Apple Expands in Austin*, Press Release, (Nov. 20, 2019) https://www.apple.com/newsroom/2019/11/apple-expands-in-austin/.

planned $1 billion, 3-million-square-foot, 133-acre campus that will house up to 15,000 employees.[4] The Apple "Austin campus" will include a 192-room hotel,[5] and provide a wide variety of engineering, research and development, operations, finance, sales and customer support jobs.[6] On information and belief, Apple also owns significant real estate in and around Austin, Texas, including for example, a 217,000 square-foot building at 320 S. Capital of Texas Hwy.[7]

28.   Apple's new Austin campus is responsible for manufacturing the Mac Pro, one of the Accused Products in this case. The Mac Pro and other Accused Products also contain components from suppliers across the state of Texas, including suppliers in this District. According to Apple's Chief Executive Officer, Tim Cook, "[b]uilding the Mac Pro, Apple's most powerful device ever, in Austin is both a point of pride and . . . deepening our close bond with the city and the talented and diverse workforce that calls it home."[8]

29.   Apple also employs individuals in this District who, on information and belief, contribute to Apple's design and sales of infringing products. For example, and without limitation, Apple employs over 100 individuals in the Austin metropolitan area on the "Siri Team" and/or with the title of SoC engineer.[9]

---

[4] *Id.*

[5] Juli Clover, *Apple Adding 192-Room Hotel to Upcoming Austin Campus*, MacRumors (May 20, 2020) https://www.macrumors.com/2020/05/20/apple-austin-campus-hotel/.

[6] Alex Caprariello, Billy Gates, *Apple's Billion-Dollar Austin Campus Nearly Finished, Move-in Date Set*, KXAN (Apr. 26, 2021), https://www.kxan.com/news/business/apples-1b-austin-campus-nearly-finished-move-in-set-for-2022/.

[7] Parimal Rohit, *Apple Buys Austin Office Building on Capital of Texas Hwy*, Austin Business Journal (Aug. 25, 2021), https://www.bizjournals.com/austin/news/2021/08/25/apple-buys-capital-ridge-austin.html.

[8] *Id.*; *see also Apple Expands in Austin*, Press Release (Nov. 20, 2019) https://www.apple.com/newsroom/2019/11/apple-expands-in-austin/.

[9] *See, e.g.* Rubinia Leal Cavazos ("Team Lead, Siri Team at Apple"), LinkedIn profile available online at https://www.linkedin.com/in/rubinialeal/; Nicole Mohrmann ("Team Lead (Siri) at Apple"), LinkedIn profile available online at https://www.linkedin.com/in/nicole-mohrmann-

30.   Apple has also sought to recruit individuals for employment within this District through its website, LinkedIn, and other websites. This includes recent postings seeking to fill hundreds of full-time positions for SoC engineers and/or Artificial Intelligence and Machine Learning Annotation Analyst on the Siri Team.[10] Most of these positions require an "On-site" presence in this District.[11]

31.   Moreover, as described below, many of the Accused Products implement Apple's infringing speech recognition technology on Intel chips. Intel is a designer and manufacturer of various processors and System on Chips. Intel is headquartered in Austin, Texas, where it employs close to 2,000 people.[12] In pertinent part, this means that on information and belief, components central to Apple's acts of infringement were designed and developed in this District, by individuals who reside and work here.

32.   On information and belief, critical documents regarding the Accused Products are also located in this District given the presence of Intel and Apple employees that live in the District and work on Apple's infringing speech recognition technology.

966ab771/; Mortuza Lokhandwala ("SOC Design Verification Engineer at Apple), LinkedIn profile available online at https://www.linkedin.com/in/murtuza-lokhandwala-a8b53785/; Vu Le ("SoC Design Engineer at Apple"), LinkedIn profile available online at https://www.linkedin.com/in/vutle/; Neil Panchal ("SoC Physical Design Manager at Apple"), LinkedIn profile available online at https://www.linkedin.com/in/neil-panchal-194a783/.

[10] *See, e.g.* Apple LinkedIn Job postings, https://www.linkedin.com/jobs/search/?current-JobId=2848276684&geoId=102748797&keywords=siri%20apple&location=Texas%2C%20United%20States; *see also* Careers at Apple: Find your perfect role (search "Austin, TX" in Location field and "SoC Architect" and/or "Siri" in search field), https://jobs.apple.com/en-us/search?search=SoC%20Architect&sort=relevance&location=austin-AST.

[11] *Id*.

[12] *See* Innovating and Investing in Texas, https://www.intel.com/content/www/us/en/corporate-responsibility/intel-in-texas.html.

ZENTIAN BREAKS NEW GROUND

33.   Mark Catchpole, Zentian's founder and one of the named inventors on the Asserted Patents, holds a B.Sc. degree in Physics from the University of Edinburgh and a M.Sc. degree in Microelectronics Systems Design from Brunel University in London. After initial stints teaching physics, Catchpole joined Phoenix VLSI and later PixelFusion as an electronics engineer working on Application Specific Integrated Circuit (ASIC) and Field Programmable Gate Array (FPGA) designs. ASICs and FPGAs are two forms of integrated circuits or chips that allow for customization of hardware after manufacturing.

34.   Most of the work by Zentian was—and still is—unpaid, as funds proved difficult to raise during the early 2000s recession following the dot-com crash. Engineers worked in exchange for share options in the company, investing their own sweat equity into the potential of a hardware system-on-chip solution for speech recognition. During this time, Zentian was able to build a partial prototype on a programmable hardware emulation platform.

35.   Guy Larri joined the Zentian team in 2004 as a technical contributor. Larri holds a B.Sc. degree in Mathematics and Computer Science, as well as an engineering degree in Electrical Engineering from the University of Sydney. Larri has worked as a VLSI (very large-scale integration) chip design engineer, designing audio processor chips, cache, memory controllers, and other microprocessors.

36.   After Larri joined, Zentian's efforts shifted to the creation of a hardware accelerator for the acoustic stage of the speech recognition process that would work alongside a separate software implementation of the search stage (presumably running on a conventional mobile processor). Once the design specifications were created, the Zentian acoustic accelerator was loaded onto a programmable gate array (a type of integrated circuit designed to be customized

after manufacturing) and paired with the Sphinx open-source speech recognition software. The Zentian accelerator demonstrated a two-times speedup of the Sphinx software.

37.   The implications of the Zentian team's discovery were vast and addressed a number of shortcomings of speech recognition systems in the art that had stymied the technology's commercial development. Speeding up the speech recognition process and reducing the overall power consumption would allow for the speech recognition system to run on a cheaper and/or lower-power processor and memory system. The Zentian accelerator would allow for the overall time the system might take to respond to spoken queries to be significantly reduced to a level more in-line with human conversations. Finally, Zentian's accelerator approach would allow for the use of larger and more sophisticated acoustic models without increasing response time, facilitating better recognition accuracy and a larger recognizable vocabulary.

38.   Buoyed by their proof of concept, Catchpole and Larri attempted to market Zentian's acoustic accelerator hardware. They spoke with representatives of leading speech recognition companies, such as Nuance, IBM, Sensory, Inc., and Samsung, as well as cellular carriers, such as Hutchison 3G UK Limited. Catchpole and Larri also attended the 2006 GSM World Congress in Barcelona, where they discussed their novel speech recognition SoC designs with an executive from Qualcomm. Unfortunately, the industry at the time remained focused on software solutions for smaller problems, such as voice-controlled phone dialing.

39.   In 2006, Zentian met with a director of Apple's Silicon Engineering Group, the group responsible for the development of multiple classes of chips used in Apple mobile products.[13] The presentation that Zentian shared with Apple identified problems with then-existent

---

[13] Apple describes its Silicon Technologies group as the team "responsible for the design and manufacture of our next generation of high-performance, power-efficient processor systems on

speech recognition systems, the promise of a dedicated hardware speech recognition accelerator, and the data demonstrating the success of Zentian's own prototype. The presentation also identified that Zentian had filed patents in the United States (and in the United Kingdom) to cover its novel hardware-based approach to speech recognition systems.

40.   Apple listened to Zentian's pitch and then ceased communication with the start-up.

ZENTIAN'S PATENTS

41.   As a result of its work in the speech recognition field, Zentian has a patent portfolio that covers a range of inventions that improve the performance and efficiency of speech recognition on-device systems, including in the areas of parallel processing and circuit configurations. Five of Zentian's patents, originating from two separate patent families, are at issue in this case.

42.   U.S. Patent No. 7,979,277 (the "'277 patent") is entitled "Speech Recognition Circuit and Method" and issued on July 12, 2011. A true and correct copy of the '277 patent is attached as **Exhibit A** to this Complaint.

43.   Under 35 U.S.C. § 154(b), the '277 patent is entitled to 1,043 days of Patent Term Adjustment, extending its expiry date until July 2028.

44.   Zentian is the owner of all rights, title, and interest in and to the '277 patent, with the full and exclusive right to file suit and enforce the '277 patent, including the right to recover damages for past infringement.

45.   The '277 patent is valid and enforceable under U.S. patent laws.

---

chip, as well as the analog IP and high-speed interface evaluation those chips need." *See* Careers at Apple: Hardware, Impossible is our favorite starting point, https://www.apple.com/careers/us/hardware.html.

46.   The '277 patent claims are directed to a patent-eligible, non-abstract idea. They address, among other things, a specific improvement to the way in which conventional software-based speech recognition systems operated on electronic devices. The '277 patent claims include an architecture for a speech recognition circuit in which data is "pipelined" between different hardware componentry. The claims specify materials (including types of processors, accelerator interfaces, etc.), structures, and configurations of the claimed speech recognition circuit designed to enhance power consumption and efficiency of the system and allow for real-time processing of more complex and robust language models.

47.   U.S. Patent No. 7,587,319 (the "'319 patent") is entitled "Speech Recognition Circuit Using Parallel Processors" and issued on September 8, 2009. A true and correct copy of the '319 patent is attached as **Exhibit B** to this Complaint.

48.   Under 35 U.S.C. § 154(b), the '319 patent is entitled to 1,137 days of Patent Term Adjustment, extending its expiry date until March 2026.

49.   Zentian is the owner of all rights, title, and interest in and to the '319 patent, with the full and exclusive right to file suit and enforce the '319 patent, including the right to recover damages for past infringement.

50.   The '319 patent is valid and enforceable under U.S. patent laws.

51.   The '319 patent claims are directed to a patent-eligible, non-abstract idea. They address, among other things, a specific improvement to the way in which conventional software-based speech recognition systems operated on electronic devices. The '319 patent claims include novel uses of a plurality of lexical tree processors connected in parallel to distribute the task of performing word recognition across processors.[14] This architecture provides for improved and

---

[14] A "lexical tree" is a data structure that comprises a model of words.

faster processing, solving one of the problems of speech recognition known in the art—the computationally intensive search process that attempts to find the most likely set of words from a given lexicon. The claims specify materials (including types of processors, accelerator interfaces, etc.), structures, and configurations designed to enhance power consumption and efficiency of the system and allow for real-time processing of more complex and robust speech models.

52.   U.S. Patent No. 10,971,140 (the "'140 patent") is entitled "Speech Recognition Circuit Using Parallel Processors" and issued on April 6, 2021. A true and correct copy of the '140 patent is attached as **Exhibit C** to this Complaint.

53.   Under 35 U.S.C. § 154(b), the '140 patent is entitled to 0 days of Patent Term Adjustment, and its expiry date is in February 2023.

54.   Zentian is the owner of all rights, title, and interest in and to the '140 patent, with the full and exclusive right to file suit and enforce the '140 patent, including the right to recover damages for past infringement.

55.   The '140 patent is valid and enforceable under U.S. patent laws.

56.   The '140 patent claims are directed to a patent-eligible, non-abstract idea. They address, among other things, a specific improvement to the way in which conventional software-based speech recognition systems operated on electronic devices. The '140 patent claims include a speech recognition circuit that has one or more clusters of processors, each of which comprises a plurality of processors and an acoustic model memory for storing acoustic model data. The disclosed and claimed circuit is further configured to generate an initial score for an audio sample, which is then used to determine whether to continue processing to determine a final score by processing a larger amount of model data than that which was processed to generate the initial score. The claims specify materials (including types of processors, accelerator interfaces, etc.),

structures, and configurations of the claimed speech recognition circuit designed to enhance power consumption and efficiency of the system and allow for real-time processing of more complex and robust language models.

57.   U.S. Patent No. 10,062,377 (the "'377 patent") is entitled "Distributed Pipelined Parallel Speech Recognition System" and issued on August 28, 2018. A true and correct copy of the '377 patent is attached as **Exhibit D** to this Complaint.

58.   Under 35 U.S.C. § 154(b), the '377 patent is entitled to 0 days of Patent Term Adjustment, and its expiry date is in September 2025.

59.   Zentian is the owner of all rights, title, and interest in and to the '377 patent, with the full and exclusive right to file suit and enforce the '377 patent, including the right to recover for damages for past infringement.

60.   The '377 patent is valid and enforceable under U.S. patent laws.

61.   The '377 patent claims are directed to a patent-eligible, non-abstract idea. They address, among other things, a specific improvement to the way in which conventional software-based speech recognition systems operated on electronic devices. The '377 patent claims include, among other things, a speech recognition system comprising a device programmed to calculate a feature vector from a digital audio stream (wherein the feature vector comprises a plurality of extracted and/or derived quantities from the stream during a defined audio time frame); a second device programmed to calculate distances indicating the similarity between the calculated feature vector and a plurality of acoustic states of an acoustic model; and a third device programmed to identify spoken words in the digital audio stream using Hidden Markov Models and/or Neural Networks. The system also includes a search stage for using the calculated distances to identify words within a lexical tree (the lexical tree comprising a model of words). The claims specify

materials (including types of processors, accelerator interfaces, etc.), structures, and configurations of the claimed speech recognition system designed to reduce power consumption and improve efficiency of the system and to allow for real time processing of more complex and robust language models.

63.   U.S. Patent No. 10,839,789 (the "'789 patent") is entitled "Speech Recognition Circuit and Method" and issued on November 17, 2020. A true and correct copy of the '789 patent is attached as **Exhibit E** to this Complaint.

63.   Under 35 U.S.C. § 154(b), the '789 patent is entitled to 0 days of Patent Term Adjustment, and its expiry date is in September 2025.

64.   Zentian is the owner of all rights, title, and interest in and to the '789 patent, with the full and exclusive right to file suit and enforce the '789 patent, including the right to recover for past infringement.

65.   The '789 patent is valid and enforceable under U.S. patent laws.

66.   The '789 patent claims are directed to a patent-eligible, non-abstract idea. They address, among other things, a specific improvement to the way in which conventional software-based speech recognition systems operated on electronic devices. The '789 patent claims include an acoustic coprocessor for processing data associated with an audio signal that includes an interface for receiving at least one feature vector, a calculating apparatus for calculating distances indicating a similarity between the feature vector and acoustic state of an acoustic model, and a second interface for sending at least one distance calculated by the calculating apparatus. The claims specify materials (including types of processors, accelerator interfaces, etc.), structures, and configurations of acoustic coprocessors designed to enhance power consumption and efficiency of the system and allow for real time processing of more complex and robust language models.

67.   For ease of reference, the '277, '377, and '789 patents are part of the same patent family. The '319 and '140 patents are part of the same patent family.

APPLE'S UNAUTHORIZED USE OF ZENTIAN'S TECHNOLOGY

68.   Apple's infringing speech recognition solutions are central to shaping the user experience of many Apple products, including in Accused Products like Apple's iPhone, iPad, Watch, HomePod smart speakers, Apple TV set top boxes, and Mac personal computers.

69.   First introduced in 2011, Siri uses speech recognition and natural language processing to enable Apple devices to play music, control smart home accessories, answer general knowledge inquiries, send messages, set clocks and timers, obtain news and weather information, and even get traffic reports and translations. According to Apple, "Siri lets you stay connected without lifting a finger" and is a "faster, easier way to do all kinds of useful things."[15]

70.   Siri fulfills more than 25 billion requests each month, according to Apple's director of Siri data science and engineering.[16] Siri is, in turn, a core feature built into almost every device Apple sells, and critical to the unique user-device interface that drives Apple's market share.

71.   Even though Siri was introduced in 2011, Apple continues to add features (*e.g.*, enabling the assistant hands free through the wake phrase "Hey Siri" in 2014 (also known as a "voice trigger") and individualized voice recognition of "Hey Siri" in 2015) and mentions Siri in product demonstrations today. Most recently, during the privacy segment of Apple's Worldwide

---

[15] *See* Siri does more than ever. Even before you ask, https://www.apple.com/siri/.
[16] *See* Ry Crist, *Siri's iOS 14 Makeover Gives You Some of Your iPhone Screen Back*, CNET (June 23, 2020), https://www.cnet.com/tech/mobile/siri-voice-assistant-ios-14-makeover-gives-some-iphone-screen-back-apple-wwdc-2020/.

Developers Conference, Apple revealed that Siri would process user requests "on device," *i.e.*, without using a network connection.[17]

72.   On information and belief, on-device voice processing is important to selling Apple's Accused Products because it allays end user privacy concerns and allegedly speeds up performance and reliability.[18] By keeping the "always on" wake phrase function on-device, users don't have to worry about Apple "listening in" on their private conversations. To the extent Apple servers receive any of the user's words or content, it is not until the wake phrase is spoken and the user is directing their speech at the Accused Product.

73.   Apple's infringement of Zentian's hardware speech recognition inventions includes among other things: (1) "Wake phrase detection" for Hey Siri; (2) Dictation of text messages, notes, and documents, with integrated applications, such as Apple's Pages and/or Microsoft Word; and (3) On-device speech recognition for Siri user requests, as announced by Apple as a privacy and performance improvement at their WorldWide Developers Conference. As described in detail below, Apple's speech recognition features would not be possible without the hardware-first architectures, methods, circuits, configurations, systems, chips and approaches engineered by Zentian in the early 2000s and claimed in the Asserted Patents.

74.   Making matters worse, Apple's infringement is willful. In early 2015, Zentian once again met with Apple, this time with Apple's Patent Acquisitions department. At the meeting,

---

[17] *See* Alex Hern, *Apple Overhauls Siri to Address Privacy Concerns and Improve Performance*, The Guardian (June 7, 2021), https://www.theguardian.com/technology/2021/jun/07/apple-over-hauls-siri-to-address-privacy-concerns-and-improve-performance.

[18] *See* Mitchell Clark, *Watch Apple's Siri Blaze Through Requests with On-device Processing*, The Verge (June 7, 2021)*,* https://www.theverge.com/2021/6/7/22523438/apple-iphone-siri-on-de-vice-processing-ios-15-wwdc-privacy (describing 2019 incident in which Apple contractors were able to listen to Siri requests on Apple servers, requiring Apple to increase transparency by making Siri recordings opt-in).

Zentian and Apple representatives discussed the two patent families at issue, including several of the then-issued Asserted Patents, such as the '277 patent and the '319 patent, as well as then-pending patent applications that kept the two patent families open for future issuances.

75.   As a result of this (and the 2006) meeting, Apple was well aware of Zentian's patented speech recognition circuits and methods when it decided to make, use, offer to sell, and/or sell the infringing Accused Products within the United States, making such acts deliberate and intentional at a minimum, and potentially willful, wanton, in bad-faith and/or consciously wrongful.

COUNT ONE: INFRINGEMENT OF U.S. PATENT NO. 7,979,277

76.   Zentian incorporates by reference and realleges the allegations in paragraphs 1-75.

77.   The '277 patent is valid and enforceable.

78.   The '277 patent names Guy Larri, Mark Catchpole, Damian Kelly Harris-Dowsett, and Timothy Brian Reynolds as inventors.

79.   Apple has directly infringed, continues to infringe, and has induced or contributed to the infringement of the '277 patent by making, using, selling, offering for sale, and/or importing into the United States, without authority or license, the Accused Products in violation of 35 U.SC. § 271(a).

80.   The Accused Products are non-limiting examples that were identified based on publicly available information. Zentian reserves the right to identify additional infringing activities, products, and services, including based on information obtained during discovery.

81.   By way of example only, the Accused Products meet all the limitations of at least independent claims 14, 15, and 16 of the '277 patent, either literally or under the doctrine of equivalents. Claim 14 is directed to a speech recognition circuit, claim 15 to a speech recognition

method, and claim 16 to a non-transitory storage medium storing processor implementable code

for controlling at least one processor to implement a speech recognition method.

82.   Claim 14 of the '277 patent recites:

A speech recognition circuit, comprising:

[1] an audio front end for calculating a feature vector from an audio signal, wherein the feature vector comprises a plurality of extracted and/or derived quantities from said audio signal during a defined audio time frame;

[2] calculating means for calculating a distance indicating the similarity between a feature vector and a predetermined acoustic state of an acoustic model; and

[3] a search stage for using said calculated distances to identify words within a lexical tree, the lexical tree comprising a model of words;

[4] wherein said audio front end, said calculating means, and said search stage are connected to each other to enable pipelined data flow.

83.   Regarding the preamble of claim 14, to the extent the preamble is determined to be

limiting, the Accused Products contain a speech recognition circuit that allows for device users to

speak the words "Hey Siri" and activate Apple's personal assistant hands-free. The speech

recognition circuit of the Accused Products further allows for dictation of text messages, notes,

and documents with integrated applications, as well as on device speech recognition for Siri user

requests, as announced by Apple recently at their Worldwide Developers Conference.[19]

84.   On information and belief, the speech recognition circuit of the Accused Products

works with cloud-based programming, natural language interpretation and other services with

---

[19] *See* Alex Hern, *Apple Overhauls Siri to Address Privacy Concerns and Improve Performance*, The Guardian (June 7, 2021)*,* https://www.theguardian.com/technology/2021/jun/07/apple-over-hauls-siri-to-address-privacy-concerns-and-improve-performance.

hardware-assisted on-device processing, creating a combination of hardware and software features.[20]

85.　　　Apple published further evidence of the speech recognition circuit of the Accused Products in October 2017. The Apple Siri Team posted an article on Apple's website entitled "Hey Siri: An On-device DNN-powered Voice Trigger for Apple's Personal Assistant."[21] That article states that a "specialized speech recognizer" incorporated on the local devices themselves "runs all the time and listens" for the words "Hey Siri," which allows for the feature to operate hands-free and respond even when the Accused Product does not have a data connection.[22]

86.　　　As to the first limitation of claim 14 of the '277 patent, the Accused Products calculate a feature vector from an audio signal using an audio front end, wherein the feature vector comprises a plurality of extracted and/or derived quantities from said audio signal during a defined audio time frame.

87.　　　In the technical paper published by Apple's Siri Team, Apple describes the Accused Products' calculation of a feature vector from an audio signal in the following way:

> The microphone in an iPhone or Apple watch turns your voice into a stream of instantaneous waveform samples, at a rate of 16000 per second. A spectrum analysis stage converts the waveform sample stream to a sequence of frames, each describing the sound spectrum of approximately 0.01 sec. About twenty of these frames at a time (0.2 sec of audio) are fed to the acoustic model, a Deep Neural Network (DNN) which converts each of these acoustic patterns into a probability distribution over a set of speech sound classes: those
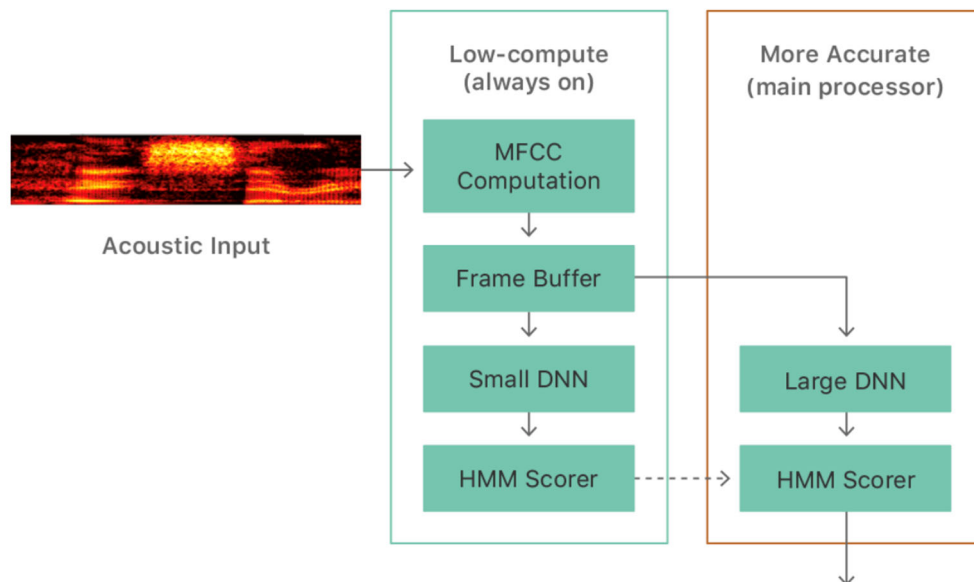
---

[20] *See, e.g.*, Christian Ziberg, *New Machine Learning Article from Apple Goes In Depth on How "Hey Siri" Does Its Magic*, iDB (Oct. 18, 2017), https://www.idownload-blog.com/2017/10/18/apple-hey-siri-machine-learning-article/.

[21] Siri Team, *Hey Siri: An On-device DNN-powered Voice Trigger for Apple's Personal Assistant* (Oct. 2017), https://machinelearning.apple.com/research/hey-siri ("*Hey Siri*"). DNN means "Deep Neural Network."

[22] *Id.*

used in the "Hey Siri" phrase, plus silence and other speech, for a total of about 20 sound classes.[23]

88.   As shown below in a figure from the Siri Team's paper, the "Always On Processor (AOP)" of the Accused Products is "a small, low-power auxiliary processor" that has access to "the microphone signal."[24] The AOP calculates feature vectors through an extraction method known as MFCC (or Mel Frequency Cepstrum Coefficient) calculations. MFCCs are a collection of coefficients that represent the short-term power spectrum of a sound.[25]
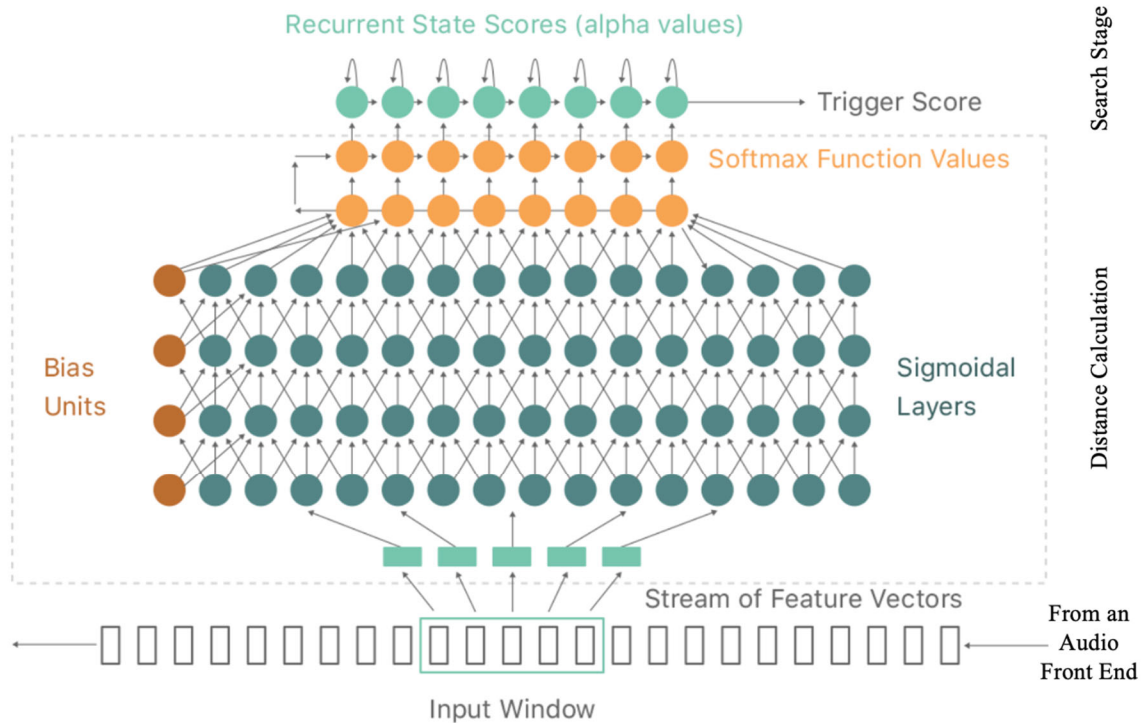

89.   The below speech recognizer diagram (annotation added in dark text to the right), also published in the Siri Team's technical paper, illustrates generally how the speech recognizer converts a spoken utterance of the words "Hey Siri" into a function performed by the Accused Product, such as waking the device. Relevant to the first limitation, the "input window" (green box

---

[23] *Id.*

[24] *Id.*

[25] *See, e.g.*, K.S. Rao and Manjunath K.E., *Speech Recognition Using Articulatory and Excitation Source Features*, Appendix A: MFCC Features at 85, SpringerBriefs in Speech Technology, https://link.springer.com/content/pdf/bbm%3A978-3-319-49220-9%2F1.pdf.

at bottom) receives a "stream of feature vectors" calculated by an audio front end (the MFCC

Computation on the AOP).[26]

90.   The Accused Products similarly contain the second limitation of claim 14—a

calculating means for calculating a distance indicating the similarity between a feature vector and

a predetermined acoustic state of an acoustic model.

91.   The acoustic model contained in the Accused Products is a Deep Neural Network

(DNN). The DNN converts the "stream of feature vectors" it receives into probability distributions

consisting of extracted and derived quantities.[27]

92.   In the Siri Team's technical paper, Apple states that the "output of the acoustic

model provides a distribution of scores over phonetic classes for every frame. A phonetic class is

---

[26] *Hey Siri* at Fig. 2 (annotations added).
[27] *Id.*

typically something like 'the first part of an /s/ preceded by a high front vowel and followed by a front vowel.'"[28] Phonetic classes of this type are otherwise known in the art as a triphones.

93.   In the same technical paper, Apple describes the DNN acoustic model (shown in the speech recognizer diagram above bounded by a dashed grey box) as follows:

> DNN consists mostly of matrix multiplications and logistic nonlinearities. Each "hidden" layer [e.g. circles connected with arrows identified as sigmoidal layers, bias units, and/or softmax function values] is an intermediate representation discovered by the DNN during its training to convert the filter bank inputs to sound classes. The final nonlinearity is essentially a Softmax function (a.k.a. a general logistic or normalized exponential).[29]

94.   The underlying mathematical equation (below) for this "distance" calculation that is the output of the acoustic model is represented by the variable $q_{i,t}$—"the log score for the phonetic class associated with the *ith* state given the acoustic pattern around time *t*"—a further indication that the acoustic model contains predetermined acoustic states.[30]

$$F_{i,t} = \max\{s_i + F_{i,t-1}, m_{i-1} + F_{i-1,t-1}\} + q_{i,t}$$

where

- $F_{i,t}$ is the accumulated score for state i of the model
- $q_{i,t}$ is the output of the acoustic model—the log score for the phonetic class associated with the *ith* state given the acoustic pattern around time t
- s_i is a cost associated with staying in state i
- m_i is a cost for moving on from state i

---

[28] *Id.*
[29] *Id.*
[30] *Id.*

95.   On information and belief, due to its power and computational requirements, the calculating means for the "Large DNN" (as compared to the "Small DNN" on the AOP) in the Accused Products is a separate processor, such as a Neural Engine with 16-core design processor in the Apple Silicon M1 Mac,[31] or alternatively, for example, in iPhone 12 and iPad Air 4th generation models that use the A14 Bionic SoC, the Graphics Processing Unit (GPU) depicted at the top of the A14 Bionic chip die image to the right.[32]


96.   The Accused Products contain the third limitation of claim 14—a search stage to identify words within a lexical tree using the calculated distances, with the lexical tree comprising a model of words.

---

[31] *See* Press Release, *Introducing M1 Pro and M1 Max: The Most Powerful Chips Apple Has Ever Built* (Oct. 18, 2021), https://www.apple.com/newsroom/2021/10/introducing-m1-pro-and-m1-max-the-most-powerful-chips-apple-has-ever-built/.

[32] *See, e.g.*, Ben Morris, *Imagination Technologies: A British Technology Success Story*, BBC (May 16, 2012), https://www.bbc.co.uk/news/business-17939014 (quoting Apple GPU provider for certain iPhone models that "Apple's popular voice command system, Siri, is driven by the GPU").

97.   In the example of the "Hey Siri" system, the speech recognizer contains a "language-specific phonetic specification of the 'Hey Siri' phrase" that was created with two English variants: one with Siri as in "serious" and two with Siri as in "Syria."[33] According to the Siri Team's technical paper on Apple's website, this model of words was built to "cope with a short break between the two words, especially as the phrase is often written with a comma: 'Hey, Siri.' Each phonetic symbol results in three speech sound classes (beginning, middle, and end) each of which is its own output from the acoustic model."[34]

98.   In the speech recognizer diagram, the search stage is depicted by the final top-layer of light green circles outside of the dotted-line box that surrounds the computationally intensive acoustic model. Combining that image with the mathematical equation above, visually shows that the output of the acoustic model, q_{it}, is fed into the search stage, which Apple describes as an "application of dynamic programming, and can be derived based on ideas about Hidden Markov Models—HMMs."[35]

99.   Further support for the existence of a search stage to identify words exists in Apple's instructions to application developers looking to incorporate the Siri technology. In the SiriKit for developers, Apple describes a "vocabulary step," consisting of a vocabulary database

---

[33] *Hey Siri*.
[34] *Id.*
[35] *Id.*

that can be adapted and added to, and that comes into play between when Siri recognizes speech from the user and creating an action that Siri can execute.[36]

100.　　　Further, the Siri Team's technical paper states that "[a]lmost all the computation in the 'Hey Siri' detector is in the acoustic model [*i.e.* Large DNN]. The temporal integration computation is relatively cheap, so we disregard it when assessing size or computational resources."[37]

101.　　　On information and belief, the temporal integration computation that occurs in the search stage occurs on either the AOP or another processor, such as one of the main application processors of the Apple system on chip (SoC). For example, there are six application processor CPU cores on the Apple A14 Bionic chip, as shown on the Tech Insights die photo (two of which are labelled "CPU1" and four of which are labelled "CPU2."[38]

102.　　　On information and belief, the high load of computing the acoustic model outputs explain why those computations are performed by the GPU or Neural Engine (depending on which SoC is in the product), where such calculations can be evaluated more efficiently and quickly, and in parallel with other workloads on the AOP and Application processors.

103.　　　As to the fourth limitation of claim 14 of the '277 patent, the Accused Products are enabled to pipeline data from an audio front end to a calculating means and then to the search stage.
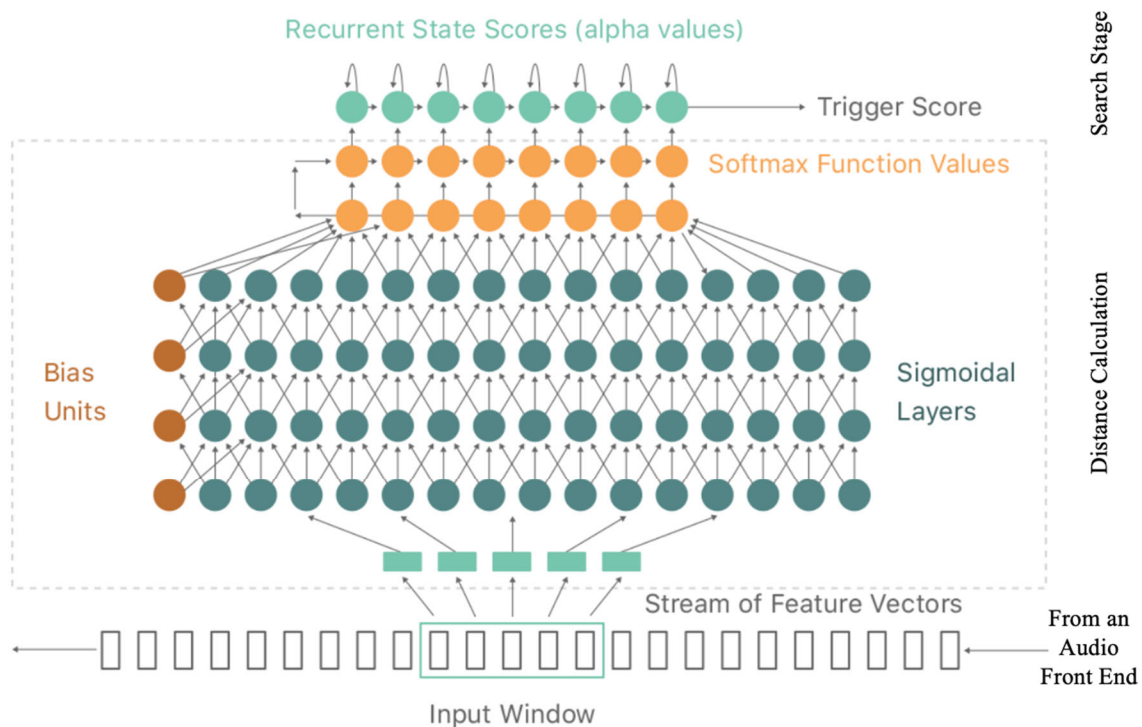
104.　　　The diagram of Apple's speech recognition system (annotated and reproduced again below) describes the flow of data in Apple's speech recognizer. On information and belief,

---

[36] *See* YML, *Learn How to Integrate Siri with Your App Thanks to Apple SiriKit* (June 29, 2016), https://ymedialabs.medium.com/learn-how-to-integrate-siri-with-your-app-thanks-to-apple-siri-kit-c6e570d302cb#1.
[37] *Hey Siri*.
[38] *Id.*

data flows in a pipelined manner from initial feature vector extraction on Apple's AOP (MFCC calculation on the AOP supplying "stream of feature vectors" to the green box labeled "Input Window" showing transmission of feature vectors), to a calculating means that computes the outputs of an acoustic model (circles and arrows within grey dotted-line box) occurring on a second processor, likely a Neural Engine or GPU, to the search stage (top line culminating in "Trigger Score") that identifies corresponding words or phrases from a language model (*i.e.* a lexical tree) that runs on either the AOP or a processor, such as the main application processors on the Apple A14 Bionic chip.[39]

105.   The above allegations detail how the Accused Products' operation and implementation of the "Hey Siri" voice trigger meets the limitations of claim 14 of the '277 patent.

---

[39] *Id.* at Fig. 2 (annotations added).

On information and belief, other speech recognition features of the Accused Products use similar technologies and implementations, are therefore similarly meet the claim limitations.

106.   Apple actively, knowingly, and intentionally induces infringement of one or more claims of the '277 patent under 35 U.S.C. § 271(b) by actively encouraging others to make, use, sell, and offer to sell in the United States, the Accused Products.

107.   For example, Apple provides directions, instruction manuals, guides, and/or other materials that encourage and facilitate infringing use by others.

108.   Apple has sold and is selling the Accused Products with the knowledge and intent that customers who buy the products will use the products for their infringing use and therefore that customers have been and are directly infringing the '277 patent.

109.   Apple contributes to the infringement of one or more claims of the '277 patent under 35 U.S.C § 271(c) by offering to sell, selling, or importing into the United States, a component of the Accused Products that constitutes a material part of the inventions, knowing the same to be especially made or especially adapted for use in an infringement of the '277 patent, and is not a staple article or commodity of commerce suitable for substantial non-infringing use.

110.   Apple was aware of the '277 patent at least from the date of this Complaint.

111.   On information and belief, Apple has been aware of the '277 patent since its issuance on July 12, 2011, and/or since Zentian representatives met with Apple's patent acquisition department in 2015 and discussed the '277 patent, among other Asserted Patents.

112.   On information and belief, Apple knew that the Accused Products infringe the '277 patent, or at a minimum believed there was a high probability that the Accused Products were covered by Zentian's patents, but willfully blinded itself to Zentian's patents and the infringing nature of the Accused Products.

113.   The foregoing allegations are based on publicly available information and a
reasonable investigation of the structure and operation of the Accused Products.

114.   Zentian reserves the right to modify this description of alleged infringement,
including, for example, on the basis of information about the Accused Products that it obtains
during discovery.

115.   The infringement details provided for claim 1 of the '277 patent are exemplary, and
Zentian intends to assert additional claims of the '277 patent beyond claim 14.

116.   Apple's infringement has damaged and continues to damage Zentian in an amount
to be determined at trial, but in no case less than a reasonable royalty for Apple's unauthorized use
of the claimed inventions set forth in the '277 patent.

117.   This is an exceptional case. Zentian is entitled to attorneys' fees and costs under 35
U.S.C § 285 as a result of the infringement of the '277 patent by Apple.

COUNT TWO: INFRINGEMENT OF U.S. PATENT NO. 7,587,319

118.   Zentian incorporates by reference and realleges the allegations in paragraphs 1-117.

119.   The '319 patent is valid and enforceable.

120.   The '319 patent names Mark Catchpole as its sole inventor.

121.   Apple has directly infringed, continues to infringe, and has induced or contributed
to the infringement of the '319 patent by making, using, selling, offering for sale, and/or importing
into the United States, without authority or license the Accused Products in violation of 35 U.SC.
§ 271(a).

122.   The Accused Products are non-limiting examples identified based on publicly
available information. Zentian reserves the right to identify additional infringing activities,
products, and services, including, for example, on the basis of information obtained during
discovery.

-29-

123.   By way of example only, the Accused Products meet all the limitations of at least

claim 46 of the '319 patent, either literally or under the doctrine of equivalents, which recites:

A speech recognition circuit, comprising:

[1] an input buffer receiving processed speech parameters:

[2] a plurality of lexical memories containing in combination complete lexical data for word recognition, each lexical memory containing part of said complete lexical data;

[3] a plurality of processors connected in parallel to said input buffer for processing the speech parameters in parallel, said processors being arranged in groups of processors, each group of processors being connected to a lexical memory:

[4] a control processor controlling each processor to process said speech parameters using partial lexical data read from a respective said lexical memory; and

[5] a results memory storing the results of the processing of the speech parameters from said processors.

124.   Regarding the preamble of claim 46, to the extent the preamble is determined to be

limiting, the Accused Products contain a speech recognition circuit that allows for device users to

speak the words "Hey Siri" and activate Apple's personal assistant hands-free. The speech

recognition circuit of the Accused Products further allows for dictation of text messages and notes

with integrated applications, as well as on device speech recognition for Siri user requests.[40]

125.   On information and belief, the speech recognition circuit of the Accused Products

works with cloud-based programming, natural language interpretation and other services with

hardware-assisted on-device processing, creating a combination of hardware and software

features.[41]

---

[40] *See* Alex Hern, *Apple Overhauls Siri to Address Privacy Concerns and Improve Performance*, The Guardian (June 7, 2021), https://www.theguardian.com/technology/2021/jun/07/apple-over-hauls-siri-to-address-privacy-concerns-and-improve-performance.
[41] *See, e.g.*, Christian Ziberg, *New Machine Learning Article from Apple Goes In Depth on How "Hey Siri" Does Its Magic*, iDB (Oct. 18, 2017), https://www.idownload-blog.com/2017/10/18/apple-hey-siri-machine-learning-article/.

126.　　　In October 2017, Apple's Siri Team published an article on Apple's website, entitled "Hey Siri: An On-device DNN-powered Voice Trigger for Apple's Personal Assistant."[42] That article states that a "specialized speech recognizer" incorporated on the local devices themselves "runs all the time and listens" for the words "Hey Siri," which allows for the feature to operate hands-free.[43] This feature also allows the Accused Product to respond to "Hey Siri" requests when a data connection is not present.
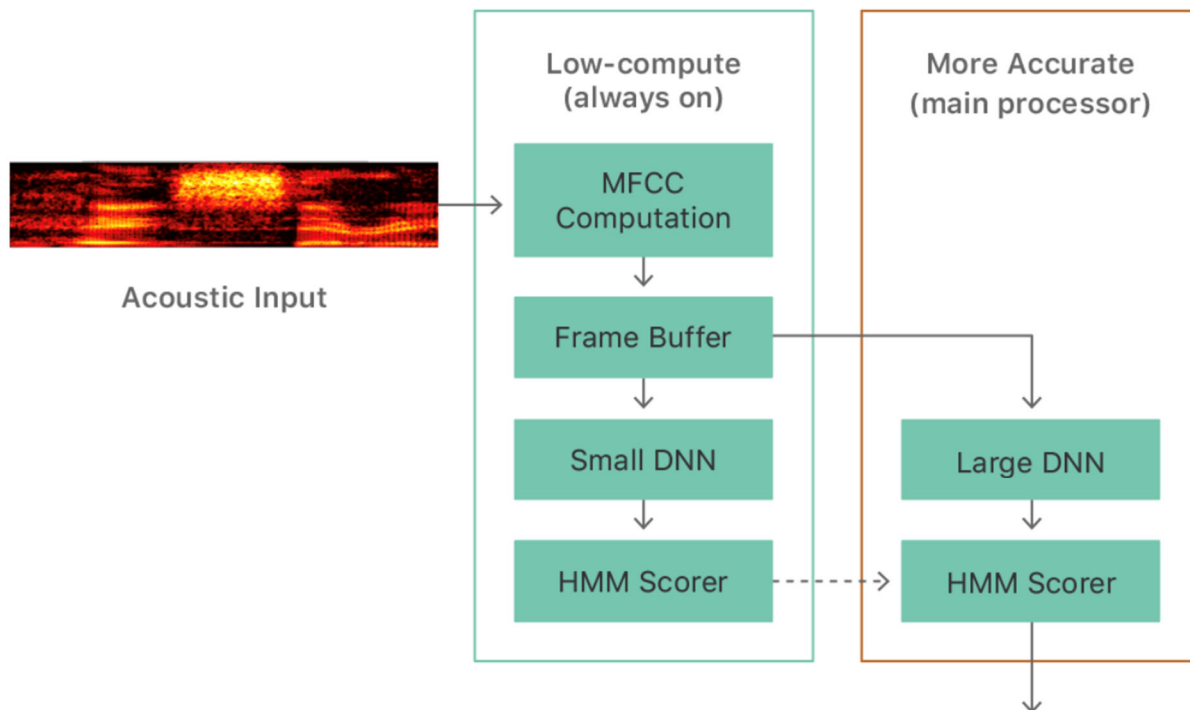
127.　　　As to the first limitation of claim 46 of the '319 patent, the Accused products contain an input buffer receiving processed speech parameters. Notably, Apple's technical paper describes a "Frame Buffer" in the AOP that stores processed speech parameters (feature vectors). The AOP receives audio input from an acoustic input, such as the microphone[44], and processes the voice audio by performing MFCC computations to produce the feature vectors that are stored in the Frame Buffer. The contents of the Frame Buffer can be received as an input to the so called "main processor," meaning that the Frame Buffer acts as an input buffer for the main processor. This architecture is shown in the figure below.[45]

---

[42] *Hey Siri*.
[43] *Id.*
[44] *Id.* ("The microphone in an iPhone or Apple Watch turns your voice into a stream of instantaneous waveform samples, at a rate of 16000 per second. A spectrum analysis stage converts the waveform sample stream to a sequence of frames, each describing the sound spectrum of approximately 0.01 sec. About twenty of these frames at a time (0.2 sec of audio) are fed to the acoustic model, a Deep Neural Network (DNN) which converts each of these acoustic patterns into a probability distribution over a set of speech sound classes.")
[45] *Id.*

128.   As to the second limitation of claim 46 of the '319 patent, the Accused Products

contain a plurality of lexical memories containing (in combination) complete lexical data for word

recognition. Notably, each Accused Product contains FLASH memory that acts as a non-volatile

memory to store data even when the device is turned off or the battery is depleted.[46] For example,

the iPhone 13 Pro A2636 256GB Model contains Kioxia 256 GB NAND Flash.[47]

129.   In addition to the FLASH memory, each Accused Product contains a cache memory

hierarchy that holds copies of data from the FLASH memory when it is accessed by processing

elements, such as the main or application processor referenced in the above figure. On information

and belief, each application processor has a private level 1 data cache and shared caches lower in

---

[46] *See, e.g.,* iPhone 13 specifications, https://www.apple.com/iphone-13/specs/ (describing FLASH memory as "Capacity" or "Storage Capacity").
[47]   *See* TechInsights, *Apple iPhone 13 Pro Teardown Briefing,* https://www.techinsights.com/blog/teardown/apple-iphone-13-pro-teardown.

the cache hierarchy, known as level 2 caches, level 3 caches, or "system caches." The FLASH memory and the caches in the cache hierarchy that lead to the Application Processors are, separately or in combination, the lexical memories containing complete lexical data for word recognition, with each containing some part of that data at any given time.

130.   On information and belief, the FLASH and cache memories store Apple's speech recognition models, which use lexical data. According to Apple, Siri's "ASR system uses a weighted finite state transducer (WFST)-based decoder, as described by Paulik."[48] Paulik is a reference to a paper published in 2015 by Apple employee Matthias Paulik, who was a senior speech research scientist at the company at the time.[49] Paulik's paper states that "[w]e (presumably referring to Apple) use a WFST based decoder that employs the difference LM principle, similar to [17]. … The decoding dictionary has 523.6K entries and the entropy pruned 4-gram language model has 16 million entries."[50] On information and belief, Apple's decoding dictionaries, with their numerous entries, contain lexical data and information.
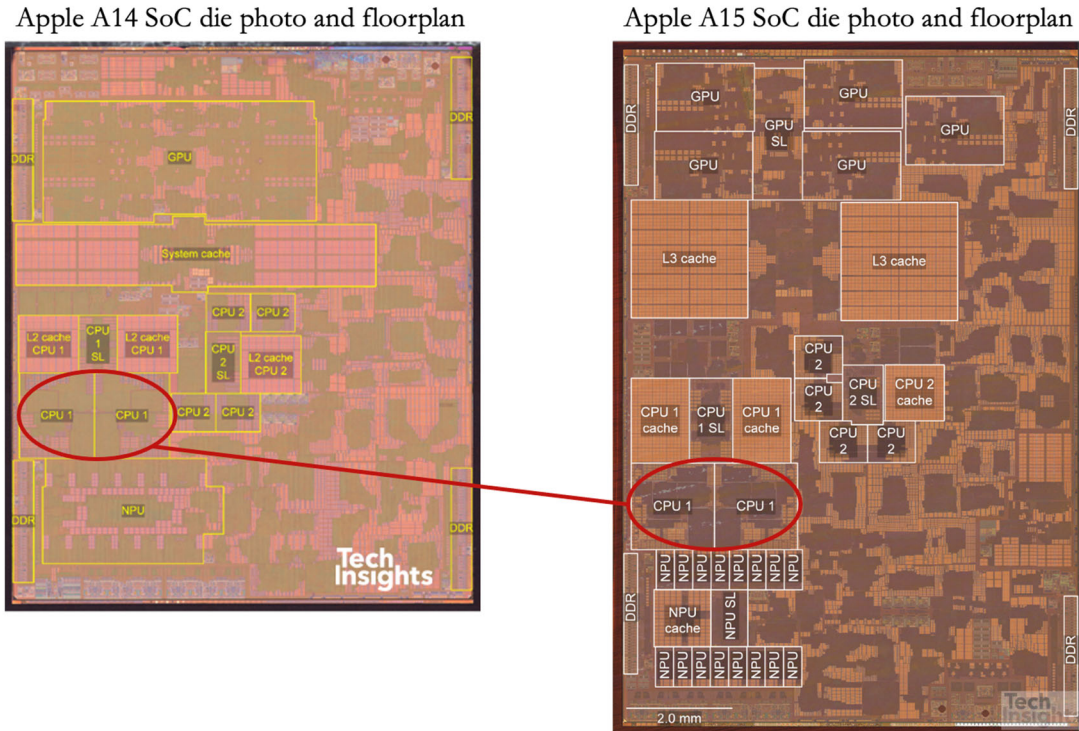
131.   As to the third limitation of claim 46 of the '319 patent, the Accused Products include more than one processor for processing speech parameters arranged in groups (with each group connected to a lexical memory). For example, the main or application processor in the Accused Products is itself arranged in groups, often referred to as "clusters." These clusters contain

---

[48] *See* Siri Speech Recognition Team, *Finding Local Destinations with Siri's Regionally Specific Language Models for Speech Recognition* (Aug. 2018), https://machinelearning.apple.com/research/regionally-specific-language-models.
[49] *See* M. Paulik, *Improvements to the Pruning Behavior of DNN Acoustic Models*, Interspeech (2015), https://www.isca-speech.org/archive_v0/interspeech_2015/papers/i15_1463.pdf .
[50] *Id.* (identifying reference [17] as Dolfing H. and Hetherington, I., *Incremental Language Models for Speech Recognition using Finite-State Transducers*, ASRU, (December 2001) Madonna di Campiglio, Trento, Italy).

multiple processors. For example, in the Apple A14 and A15 SoC as seen in the die photos below, Cluster 1 contains two application processors labelled CPU1.[51]

Apple A14 SoC die photo and floorplan

Apple A15 SoC die photo and floorplan

132.   Moreover, as further shown above, each of the processors within a cluster is connected to cache memory, including a private cache, a level 2 cluster level shared cache, and a level 3 system cache ("L3 cache" in A15 SoC). On information and belief, these caches along with the Accused Products' FLASH memory are sources of lexical memory because they store some part (or all) of the lexical data for word recognition.

133.   On information and belief, these processors and clusters are connected in parallel to the input buffer. As described and shown above, during the speech recognition process, the

---

[51] *See* Tech Insights, *Two New Apple SoCs, Two Market Events: Apple A14 and M1*, https://www.techinsights.com/blog/two-new-apple-socs-two-market-events-apple-a14-and-m1 (A14 SoC die photo); *Apple iPhone 13 Pro Teardown Briefing*, https://www.techinsights.com/blog/teardown/apple-iphone-13-pro-teardown (A15 SoC die photo).

contents of the "Frame Buffer" in the AOP can be received as an input to the so called "main

processor" (likely one or more of the clusters of processors described above, such as CPU1 and/or

CPU2), meaning that the Frame Buffer acts as an input buffer for the main processor.

134.   The Accused Products also contain the fourth limitation of claim 46 of the '319

patent, which is a control processor that controls each processor that is processing speech

parameters using lexical data read from stored lexical memory. In a technical paper published in

2017, Apple's Siri Team described one function of Apple's speech recognition circuit—the "Hey

Siri" wake phrase detection in the following manner:

> To avoid running the main processor all day just to listen for the trigger phrase, the
> iPhone's Always On Processor (AOP) (a small, low-power auxiliary processor, that
> is the embedded Motion Coprocessor) has access to the microphone signal (on 6S
> and later). We use a small proportion of the AOP's limited processing power to run
> a detector with a small version of the acoustic model (DNN). When the score
> exceeds a threshold the motion coprocessor wakes up the main processor, which
> analyzes the signal using a larger DNN.[52]

135.   On information and belief, the AOP controls the main or application processor, as

described in the paragraph above. As further described, the main or application processor processes

speech parameters using lexical data—the data within the large DNN. And that lexical data is

stored in the FLASH and/or cache memory structures connected to and shared by those processors.

136.   Finally, the fifth limitation of claim 46 of the '319 is also present in the Accused

Products, as each has a results memory storing the results of processed speech parameters. Notably,

each of the Accused Products contain SDRAM (Synchronous Dynamic Random Access

Memory).[53]

---

[52] *Hey Siri*.
[53] *See, e.g.,*   TechInsights, *Apple   iPhone   13   Pro   Teardown   Briefing*,
https://www.techinsights.com/blog/teardown/apple-iphone-13-pro-teardown.

137.   On information and belief, SDRAM is commonly the main working memory of the computation system in electronic devices like the Accused Products and would store the results of the processed speech parameters from the AOP and main processors.

138.   Apple actively, knowingly, and intentionally induces infringement of one or more claims of the '319 patent under 35 U.S.C. § 271(b) by actively encouraging others to make, use, sell, and offer to sell in the United States, the Accused Products.

139.   For example, Apple provides directions, instruction manuals, guides, and/or other materials that encourage and facilitate infringing use by others.

140.   Apple has sold and is selling the Accused Products with the knowledge and intent that customers who buy the products will use the products for their infringing use and therefore that customers have been and are directly infringing the '319 patent.

141.   Apple contributes to the infringement of one or more claims of the '319 patent under 35 U.S.C § 271(c) by offering to sell, selling, or importing into the United States, a component of the Accused Products that constitutes a material part of the inventions, knowing the same to be especially made or especially adapted for use in an infringement of the '319 patent, and is not a staple article or commodity of commerce suitable for substantial noninfringing use.

142.   Apple was aware of the '319 patent at least from the date of this Complaint.

143.   On information and belief, Apple has been aware of the '319 patent since its issuance on September 8, 2009, and/or since Zentian representatives met with Apple's patent acquisition department in 2015 and discussed the '319 patent, among other Asserted Patents.

144.   On information and belief, Apple knew that the Accused Products infringe the '319 patent, or at a minimum believed there was a high probability that the Accused Products were

covered by Zentian's patents, but willfully blinded itself to Zentian's patents and the infringing nature of the Accused Products.

145.   The foregoing allegations are based on publicly available information and a reasonable investigation of the structure and operation of the Accused Products.

146.   Zentian reserves the right to modify this description of alleged infringement, including, for example, on the basis of information about the Accused Products that it obtains during discovery.

147.   The infringement details provided for claim 1 of the '319 patent are exemplary, and Zentian intends to assert additional claims of the '319 patent beyond claim 46.

148.   Apple's infringement has damaged and continues to damage Zentian in an amount to be determined at trial, but in no case less than a reasonable royalty for Apple's unauthorized use of the claimed inventions set forth in the '319 patent.

149.   This is an exceptional case. Zentian is entitled to attorneys' fees and costs under 35 U.S.C § 285 as a result of the infringement of the '319 patent by Apple.

COUNT THREE: INFRINGEMENT OF U.S. PATENT NO. 10,971,140

150.   Zentian incorporates by reference and realleges the allegations in paragraphs 1-149.

151.   The '140 patent is valid and enforceable.

152.   The '140 patent names Mark Catchpole as its sole inventor.

153.   Apple has directly infringed, continues to infringe, and has induced or contributed to the infringement of the '140 patent by making, using, selling, offering for sale, and/or importing into the United States, without authority or license the Accused Products in violation of 35 U.SC. § 271(a).

154.   The Accused Products are non-limiting examples identified based on publicly available information. Zentian reserves the right to identify additional infringing activities,

products, and services, including, for example, on the basis of information obtained during discovery.

155.	By way of example only, the Accused Products meet all the limitations of at least claim 1 of the '140 patent, either literally or under the doctrine of equivalents, which recites:

A speech recognition circuit, comprising:

[1] one or more cluster of processors, each of the one or more cluster of processors comprising:

[2] a plurality of processors; and

[3] an acoustic model memory storing acoustic model data, wherein each of the plurality of processors is configured to compute a probability using the acoustic model data in the acoustic model memory, wherein:

[4] the speech recognition circuit is configured to generate an initial score for an audio sample; and

[5] the initial score is used to determine whether to continue processing to determine a final score via processing a larger amount of model data than that was processed to generate the initial score.

156.	Regarding the preamble of claim 1, to the extent the preamble is determined to be limiting, the Accused Products contain a speech recognition circuit that allows for device users to speak the words "Hey Siri" and activate Apple's personal assistant hands-free. The speech recognition circuit of the Accused Products further allows for dictation of text messages and notes with integrated applications, as well as on device speech recognition for Siri user requests.[54]

157.	On information and belief, the speech recognition circuit of the Accused Products works with cloud-based programming, natural language interpretation and other services with

---

[54] *See* Alex Hern, *Apple Overhauls Siri to Address Privacy Concerns and Improve Performance*, The Guardian (June 7, 2021), https://www.theguardian.com/technology/2021/jun/07/apple-over-hauls-siri-to-address-privacy-concerns-and-improve-performance.

hardware-assisted on device processing, creating a combination of hardware and software features.[55]

158.   In October 2017, Apple's Siri Team published an article on Apple's website, entitled "Hey Siri: An On-device DNN-powered Voice Trigger for Apple's Personal Assistant."[56] That article states that a "specialized speech recognizer" incorporated on the local devices themselves "runs all the time and listens" for the words "Hey Siri," which allows for the feature to operate hands-free and respond even when the Accused Product does not have a data connection.[57]

159.   As to the first limitation of claim 1 of the '140 patent, the Accused Products contain one or more clusters of processors. For example, the iPhone 13 and sixth generation iPad Mini contain an Apple A15 Bionic System on Chip (SoC). The A15 chip contains two "high performance" cores called Avalanche and four "energy efficient" cores called Blizzard. In addition, the A15 chip contains both Apple's Neural Engine and Apples Graphics Processor Unit (GPU).[58] Each of these engines and/or processors contain within them multiple "cores" and/or processors, such that they can be described, either individually or collectively, as a "clusters of processors."[59]

---

[55] *See, e.g.*, Christian Ziberg, *New Machine Learning Article from Apple Goes In Depth on How "Hey Siri" Does Its Magic*, iDB (Oct. 18, 2017), https://www.idownload-blog.com/2017/10/18/apple-hey-siri-machine-learning-article/.
[56] *Hey Siri*.
[57] *Id.*
[58] *The New Apple A15 Bionic Chip*, Mac O'Clock (Oct. 12, 2021), https://medi-um.com/macoclock/the-new-apple-a15-bionic-chip-dc6859d2c90e.
[59] *Id.*

160.   These clusters of processors that comprise the claimed speech recognition circuit also were touted by Apple as key features of the A15 chip as seen in the specification diagram below highlighting the 30% faster GPU, faster Neural Engine, and faster CPU.

161.   The Accused Products further contain the second limitation of claim 1 of the '140 patent: notably, each cluster of processors comprises a plurality of processors. Notably, Apple's Neural Engine within the Apple A15 SoC contains 16 processor cores and performs more than 15 trillion operations per second.[60] The GPU within the A15 SoC contains 4 or 5 processor cores, depending on the model (5 for the iPhone Pro and iPhone 13 Pro Max).[61] And on information and belief, Apple's AON processor contains more than one processor within it.

---

[60] *See* Stephen Shankland, *Apple's A15 Bionic Chip Powers iPhone 13 with 15 Billion Transistors, New Graphics and AI*, CNET (Sept. 14, 2021), https://www.cnet.com/tech/mobile/apples-a15-bi-onic-chip-powers-iphone-13-with-15-billion-transistors-new-graphics-and-ai/.

[61] *Id.*; *see also* Filipe Espósito, *iPhone 13 Pro's A15 Bionic Chip Has More Powerful GPU than Regular iPhone 13*, 9 to 5 Mac (Sept. 14, 2021), https://9to5mac.com/2021/09/14/iphone-13-pros-a15-bionic-chip-has-more-powerful-gpu-than-regular-iphone-13/; *see also* Ben Morris, *Imagina-tion Technologies: A British Technology Success Story*, BBC News (May 16, 2012), https://www.bbc.co.uk/news/business-17939014 (noting that Siri "is driven by the GPU").

162.   The Accused Products further contain the third limitation of claim 1 of the '140

patent as well—an acoustic model memory storing acoustic model data, wherein each of the

plurality of processors is configured to calculate a probability using the acoustic model data in the

acoustic model memory. Among other processors, Apple's Neural Engine is used for speech

recognition. For example, in November 2020, Apple stated that its M1 SoC contained in the fifth

generation iPad Pro, 24-inch iMac (M1), 2020 13-inch MacBrook Pro, MacBook Air, and Mac

mini, "features our most advanced Neural Engine with a 16-core design. . . . The entire M1 chip is

designed to excel at machine learning, and tasks like . . . voice recognition . . . will have a level of

performance never seen before on the Mac."[62]

163.   On information and belief, Apple's Neural Engine contains an acoustic model

memory storing acoustic model data in the form of on-chip SRAM, DRAM, or cache memory, as

part of the M1 chip's unified memory architecture that makes any acoustic model memory

available for the Neural Engine, GPU, and/or CPU.[63]

---

[62] *See* Apple   November   10,   2020   Event   at   12   minutes   10   seconds,
https://www.youtube.com/watch?v=5AwdkGKmZ0I (last visited Jan. 3, 2022).
[63] *See* U.S. Patent Application Publication 2019/0340491 Figures 2 ("[t]he neural engine 314
receives   the   input   data   322,   performs   multiply-accumulate   operations   (e.g.,   convolution
operations) on the input data 322 based on stored kernel data")

164.   Moreover, the Accused Products contain the fourth and fifth limitations that the

speech recognition circuit is configured to generate an initial score for an audio sample and that

the initial score is used to determine whether to continue processing to determine a final score via

processing a larger amount of model data than what was processed to generate the initial score.

165.   In a technical paper published in 2017, Apple's Siri Team described one function

of Apple's speech recognition circuit—the "Hey Siri" wake phrase detection in the following

manner:

> To avoid running the main processor all day just to listen for the trigger phrase, the
> iPhone's Always On Processor (AOP) (a small, low-power auxiliary processor, that
> is the embedded Motion Coprocessor) has access to the microphone signal (on 6S
> and later). We use a small proportion of the AOP's limited processing power to run
> a detector with a small version of the acoustic model (DNN). When the score
> exceeds a threshold the motion coprocessor wakes up the main processor, which
> analyzes the signal using a larger DNN. In the first versions with AOP support, the
> first detector used a DNN with 5 layers of 32 hidden units and the second detector
> had 5 layers of 192 hidden units.[64]

---

[64] *Hey Siri*.

166.   In turn, the "initial score" computed by a the AOP and a small DNN containing a smaller amount of data (5 layers of 32 hidden units) is used to determine whether to continue processing to determine a final score computed by a "larger DNN" that contains a larger amount of model data than what was processed to generate the initial score (5 layers of 192 hidden units).

167.   Apple actively, knowingly, and intentionally induces infringement of one or more claims of the '140 patent under 35 U.S.C. § 271(b) by actively encouraging others to make, use, sell, and offer to sell in the United States, the Accused Products.

168.   For example, Apple provides directions, instruction manuals, guides, and/or other materials that encourage and facilitate infringing use by others.

169.   Apple has sold and is selling the Accused Products with the knowledge and intent that customers who buy the products will use the products for their infringing use and therefore that customers have been and are directly infringing the '140 patent.

170.   Apple contributes to the infringement of one or more claims of the '140 patent under 35 U.S.C § 271(c) by offering to sell, selling, or importing into the United States, a component of the Accused Products that constitutes a material part of the inventions, knowing the same to be especially made or especially adapted for use in an infringement of the '140 patent, and is not a staple article or commodity of commerce suitable for substantial noninfringing use.

171.   Apple was aware of the '140 patent at least from the date of this Complaint.

172.   On information and belief, Apple has been aware of the '140 patent since its issuance on April 6, 2021.

173.   On information and belief, Apple knew that the Accused Products infringe the '140 patent, or at a minimum believed there was a high probability that the Accused Products were

covered by Zentian's patents, but willfully blinded itself to Zentian's patents and the infringing nature of the Accused Products.

174.   The foregoing allegations are based on publicly available information and a reasonable investigation of the structure and operation of the Accused Products.

175.   Zentian reserves the right to modify this description of alleged infringement, including, for example, on the basis of information about the Accused Products that it obtains during discovery.

176.   The infringement details provided for claim 1 of the '140 patent are exemplary, and Zentian intends to assert additional claims of the '140 patent beyond claim 1.

177.   Apple's infringement has damaged and continues to damage Zentian in an amount to be determined at trial, but in no case less than a reasonable royalty for Apple's unauthorized use of the claimed inventions set forth in the '140 patent.

178.   This is an exceptional case. Zentian is entitled to attorneys' fees and costs under 35 U.S.C § 285 as a result of the infringement of the '140 patent by Apple.

COUNT FOUR: INFRINGEMENT OF U.S. PATENT NO. 10,062,377

179.   Zentian incorporates by reference and realleges the allegations in paragraphs 1-178.

180.   The '377 patent is valid and enforceable.

181.   The '377 patent names Guy Larri, Mark Catchpole, Damian Kelly Harris-Dowsett, and Timothy Brian Reynolds as inventors.

182.   Apple has directly infringed, continues to infringe, and has induced or contributed to the infringement of the '377 patent by making, using, selling, offering for sale, and/or importing into the United States, without authority or license the Accused Products in violation of 35 U.SC. § 271(a).

183.   The Accused Products are non-limiting examples identified based on publicly available information. Zentian reserves the right to identify additional infringing activities, products, and services, including, for example, on the basis of information obtained during discovery.

184.   By way of example only, the Accused Products meet all the limitations of at least independent claim 1 of the '377 patent, and dependent claims 2-6, either literally or under the doctrine of equivalents.

185.   Claim 1 of the '377 patent recites:

A speech recognition system, comprising:

[1] a first programmable device programmed to calculate a feature vector from a digital audio stream, wherein the feature vector comprises a plurality of extracted and/or derived quantities from said digital audio stream during a defined audio time frame;

[2] a second programmable device programmed to calculate distances indicating the similarity between a feature vector and a plurality of acoustic states of an acoustic model wherein said feature vector is received by the second programmable device after it is calculated by the first programmable device; and

[3] a third programmable device programmed to identify spoken words in said digital audio stream using Hidden Markov Model and/or Neural Networks wherein said word identification uses one or more distances that were calculated by the second programmable device, wherein said identification of spoken words uses one or more distances calculated from a first feature vector; and

[4] a search stage for using the calculated distances to identify words within a lexical tree, the lexical tree comprising a model of words.

186.   Regarding the preamble of claim 1, to the extent the preamble is determined to be limiting, the Accused Products contain a speech recognition system that allows for device users to speak the words "Hey Siri" and activate Apple's personal assistant hands-free. The speech

recognition system of the Accused Products further allows for dictation of text messages and notes with integrated applications, as well as on device speech recognition for Siri user requests.[65]

187.   On information and belief, the speech recognition system of the Accused Products works with cloud-based programming, natural language interpretation and other services with hardware-assisted on device processing, creating a combination of hardware and software features.[66]

188.   In October 2017, Apple's Siri Team published an article on Apple's website, entitled "Hey Siri: An On-device DNN-powered Voice Trigger for Apple's Personal Assistant."[67] That article states that a "specialized speech recognizer" incorporated on the local devices themselves "runs all the time and listens" for the words "Hey Siri," which allows for the feature to operate hands-free and respond even when the Accused Product does not have a data connection.[68]

189.   As to the first limitation of claim 1 of the '377 patent, the Accused Products contain a first programmable device programmed to calculate a feature vector from a digital audio stream, wherein the feature vector comprises a plurality of extracted and/or derived quantities from said digital audio stream during a defined audio time frame.

190.   As shown below in a figure published by Apple's Siri Team, the "Always On Processor (AOP)" of the Accused Products is a programmable device, described as "a small, low-

---

[65] *See* Alex Hern, *Apple Overhauls Siri to Address Privacy Concerns and Improve Performance*, The Guardian (June 7, 2021), https://www.theguardian.com/technology/2021/jun/07/apple-over-hauls-siri-to-address-privacy-concerns-and-improve-performance.

[66] *See, e.g.*, Christian Ziberg, *New Machine Learning Article from Apple Goes In Depth on How "Hey Siri" Does Its Magic*, iDB (Oct. 18, 2017), https://www.idownload-blog.com/2017/10/18/apple-hey-siri-machine-learning-article/.

[67] *Hey Siri*.

[68] *Id.*

power auxiliary processor" that has access to "the microphone signal."[69] The AOP calculates feature vectors through an extraction method known as MFCC (or Mel Frequency Cepstrum Coefficient) calculations. MFCCs are a collection of coefficients that represent the short-term power spectrum of a sound.[70]

191.   Apple's Siri Team further describes the Accused Products' calculation of a feature vector from an audio signal in the following way:
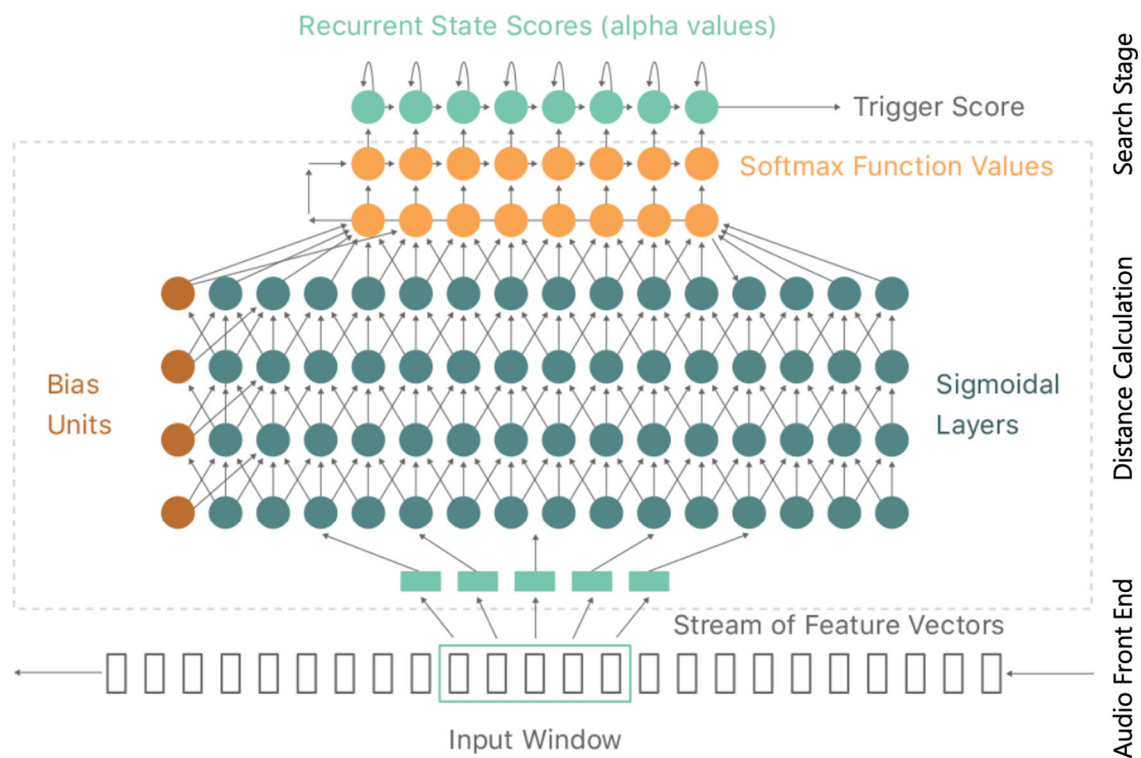
> The microphone in an iPhone or Apple watch turns your voice into a stream of instantaneous waveform samples, at a rate of 16000 per second. A spectrum analysis stage converts the waveform sample stream to a sequence of frames, each describing the sound spectrum of approximately 0.01 sec.[71]

---

[69] *Id.*

[70] *See, e.g.*, K.S. Rao and Manjunath K.E., *Speech Recognition Using Articulatory and Excitation Source Features,* Appendix A: MFCC Features at 85, SpringerBriefs in Speech Technology, https://link.springer.com/content/pdf/bbm%3A978-3-319-49220-9%2F1.pdf.

[71] *Hey Siri*.

192.   The below speech recognizer diagram (annotation added in dark text on the right), also published in the Siri Team's technical paper, illustrates generally how the speech recognition system of the Accused Products converts a spoken utterance of the words "Hey Siri" into a function, such as waking the device. Relevant to the first limitation, the "input window" (green box at bottom) receives a "stream of feature vectors" calculated by an audio front end (the MFCC Computation on the AOP).[72]

193.   The Accused Products contain the second limitation of claim 1 of the '377 patent: notably, they contain a second programmable device programmed to calculate distances (*i.e.*, probabilities) indicating the similarity between the feature vector received from the first programable device and a plurality of acoustic states of an acoustic model.

---

[72] *Id.* at Figure 2 (annotations added).

194.   As shown above in the figure published by Apple's Siri Team, the "More Accurate (main processor)" of the Accused Products is a second programmable device designed to calculate probabilities that indicate the similarity between the feature vector received from the first programable device and a plurality of acoustic states of an acoustic model.

195.   The acoustic model contained in the Accused Products is a Deep Neural Network (DNN). The DNN converts the 0.2 seconds of audio (approximately twenty frames of 0.01 seconds) it receives into probability distributions consisting of extracted and derived quantities, like the 20 speech sound classes used in "Hey Siri" phrase, plus silence and other speech.[73]

196.   In the Siri Team's technical paper, Apple states that the "output of the acoustic model provides a distribution of scores over phonetic classes for every frame. A phonetic class is typically something like 'the first part of an /s/ preceded by a high front vowel and followed by a front vowel.'" Phonetic classes of this type are otherwise known in the art as a triphones.[74]

197.   Elsewhere in the same paper, Apple describes the DNN acoustic model (shown in the figure above bounded by a dashed grey box receiving feature vectors within the "input window") as follows:

> "DNN consists mostly of matrix multiplications and logistic nonlinearities. Each 'hidden' layer [e.g. circles connected with arrows identified as sigmoidal layers, bias units, and/or softmax function values] is an intermediate representation discovered by the DNN during its training to convert the filter bank inputs to sound classes. The final nonlinearity is essentially a Softmax function (a.k.a. a general logistic or normalized exponential)."[75]

198.   The underlying mathematical equation (below) for this "distance" calculation that is the output of the acoustic model is represented by the variable $q_{i,t}$—"the log score for the

---

[73] *Id.*
[74] *Id.*
[75] *Id.*

phonetic class associated with the *ith* state given the acoustic pattern around time *t*."[76] In turn, the calculation is a mathematical comparison of the similarity between the feature vector received from the AOP and acoustic states of the DNN.

$$F_{i,t} = \max\{s_i + F_{i,t-1}, m_{i-1} + F_{i-1,t-1}\} + q_{i,t}$$

where

- $F_{i,t}$ is the accumulated score for state *i* of the model
- $q_{i,t}$ is the output of the acoustic model—the log score for the phonetic class associated with the *ith* state given the acoustic pattern around time *t*
- s_i is a cost associated with staying in state *i*
- m_i is a cost for moving on from state *i*

199. On information and belief, due to its power and computational requirements, the "Large DNN" (as compared to the "Small DNN" on the AOP) in the Accused Products runs on a separate (second) processor, which is a programmable device such as in the Apple Silicon M1 Mac, a Neural Engine with 16-core design processor.[77] Or alternatively, for example, in iPhone 12 and iPad Air 4[th] generation models that use the A14 Bionic system on chip, the Graphics Processing Unit (GPU) depicted at the top of the A14 Bionic chip die image to the right.[78]

200. The Accused Products contain the third limitation of claim 1 of the '377 patent as well— a third programmable device programmed to identify spoken words in a digital audio

---

[76] *Id.*

[77] *See* Press Release, *Introducing M1 Pro and M1 Max: The Most Powerful Chips Apple Has Ever Built* (Oct. 18, 2021), https://www.apple.com/newsroom/2021/10/introducing-m1-pro-and-m1-max-the-most-powerful-chips-apple-has-ever-built/. In November 2020, Apple stated that its M1 SoC "features our most advanced Neural Engine with a 16-core design. The entire M1 chip is designed to excel at machine learning, and tasks like . . . voice recognition . . . will have a level of performance never before seen on the Mac." Apple November 10, 2020 Event at 12 minutes 10 seconds, https://www.youtube.com/watch?v=5AwdkGKmZ0I.

[78] *See, e.g.*, Ben Morris, *Imagination Technologies: A British technology success story,* BBC News (May 16, 2012), https://www.bbc.co.uk/news/business-17939014 (quoting Apple GPU provider for certain iPhone models that "Apple's popular voice command system, Siri, is driven by the GPU").

stream using either a Hidden Markov Model and/or a Neural Network. The identification of spoken words by the third programmable device uses one or more of the distances calculated by the second programmable device, which themselves were calculated from the feature vectors calculated by the first programmable device.

201.   On information and belief, the temporal integration computation that occurs in the final stage of the speech recognition process (the search stage) identifies spoken words from the digital audio stream and occurs on a third programmable device, likely another processor, such as one of the main application processors. For example, there are six application processor CPU cores on the Apple A14 Bionic chip, as shown on the die photo above (two of which are labelled "CPU1" and four of which are labelled "CPU2").[79]

202.   In the speech recognizer diagram, this search stage is depicted by the final top-layer of light green circles outside of the dotted-line box that surrounds the computationally intensive acoustic model that functions to calculate the "distances" described above. Combining that image with the mathematical equation above shows that the output of the acoustic model, q_{it}, is fed into this last stage, which Apple describes as an "application of dynamic programming, and can be derived based on ideas about Hidden Markov Models—HMMs."[80]

---

[79] *Hey Siri*.
[80] *Id.*

203.   Moreover, the Accused Products contain the fourth limitation of claim 1 of the '377

patent: namely, they include a search stage for using the calculated distances to identify words

within a lexical tree, the lexical tree comprising a model of words.

204.   The "Hey Siri" system contains a "language-specific phonetic specification of the

'Hey Siri' phrase" that was created with two English variants: one with Siri as in "serious" and

two with Siri as in "Syria."[81] According to the Siri Team's technical paper on Apple's website,

this model of words (i.e. the lexical tree) was built to "cope with a short break between the two

words, especially as the phrase is often written with a comma: 'Hey, Siri.' Each phonetic symbol

results in three speech sound classes (beginning, middle and end) each of which as its own output

from the acoustic model."[82]

205.   Apple's instructions to application developers looking to incorporate the Siri

technology provides further support for the existence of a search stage to identify words within a

lexical tree. In the SiriKit for developers, Apple describes a "vocabulary step," consisting of a

vocabulary database that can be adapted and added to, and that comes into play between when Siri

recognizes speech from the user and creating an action that Siri can execute.[83]

206.   Apple actively, knowingly, and intentionally induces infringement of one or more

claims of the '377 patent under 35 U.S.C. § 271(b) by actively encouraging others to make, use,

sell, and offer to sell in the United States, the Accused Products.

207.   For example, Apple provides directions, instruction manuals, guides, and/or other

materials that encourage and facilitate infringing use by others.

---

[81] *Hey Siri*.
[82] *Id*.
[83] *See* YML Innovation Lab, *Learn How to Integrate Siri with Your App Thanks to Apple SiriKit*
(June 29, 2016), https://ymedialabs.medium.com/learn-how-to-integrate-siri-with-your-app-
thanks-to-apple-sirikit-c6e570d302cb#1.

208.   Apple has sold and is selling the Accused Products with the knowledge and intent that customers who buy the products will use the products for their infringing use and therefore that customers have been and are directly infringing the '377 patent.

209.   Apple contributes to the infringement of one or more claims of the '377 patent under 35 U.S.C § 271(c) by offering to sell, selling, or importing into the United States, a component of the Accused Products that constitutes a material part of the inventions, knowing the same to be especially made or especially adapted for use in an infringement of the '377 patent, and is not a staple article or commodity of commerce suitable for substantial noninfringing use.

210.   Apple was aware of the '377 patent at least from the date of this Complaint.

211.   On information and belief, Apple has been aware of the '377 patent since its issuance on August 28, 2018.

212.   On information and belief, Apple knew that the Accused Products infringe the '377 patent, or at a minimum believed there was a high probability that the Accused Products were covered by Zentian's patents, but willfully blinded itself to Zentian's patents and the infringing nature of the Accused Products.

213.   The foregoing allegations are based on publicly available information and a reasonable investigation of the structure and operation of the Accused Products.

214.   Zentian reserves the right to modify this description of alleged infringement, including, for example, on the basis of information about the Accused Products that it obtains during discovery.

215.   The infringement details provided for claim 1 of the '377 patent are exemplary, and Zentian intends to assert additional claims of the '377 patent beyond claim 1.

216.   Apple's infringement has damaged and continues to damage Zentian in an amount to be determined at trial, but in no case less than a reasonable royalty for Apple's unauthorized use of the claimed invention set forth in the '377 patent.

217.   This is an exceptional case. Zentian is entitled to attorneys' fees and costs under 35 U.S.C § 285 as a result of the infringement of the '377 patent by Apple.

COUNT FIVE: INFRINGEMENT OF U.S. PATENT NO. 10,839,789

218.   Zentian incorporates by reference and realleges the allegations in paragraphs 1-217.

219.   The '789 patent is valid and enforceable.

220.   The '789 patent names Guy Larri, Mark Catchpole, Damian Kelly Harris-Dowsett, and Timothy Brian Reynolds as inventors.

221.   Apple has directly infringed, continues to infringe, and has induced or contributed to the infringement of the '789 patent by making, using, selling, offering for sale, and/or importing into the United States, without authority or license the Accused Products in violation of 35 U.SC. § 271(a).

222.   The Accused Products are non-limiting examples identified based on publicly available information. Zentian reserves the right to identify additional infringing activities, products, and services, including, for example, on the basis of information obtained during discovery.

223.   By way of example only, the Accused Products meet all the limitations of at least claim 10 of the '789 patent, either literally or under the doctrine of equivalents, which recites:

An acoustic coprocessor, comprising:

[1] a first interface for receiving at least one feature vector;

[2] a calculating apparatus for calculating distances indicating a similarity between the at least one feature vector and at least one acoustic state of an acoustic model read from an acoustic model memory; and

[3] a second interface for sending at least one distance calculated by the calculating apparatus.

224.   Regarding the preamble of claim 10, to the extent the preamble is determined to be limiting, the Accused Products contain an acoustic coprocessor. For example, the 2020 MacBook Pro 13-Inch Core-i5 2.0GHz laptop computer contains an Intel Gaussian & Neural Accelerator (GNA). According to Intel, the Intel GNA "delivers a dedicated engine for background workloads such as voice processing and noise suppression at ultra-low power, for maximum battery life."[84]

225.   The Intel GNA is located on the "Ice Lake" Core-i5 processor on the 2020 MacBook Pro 2.0GHz, as marketed by Intel upon its release in 2019.[85]
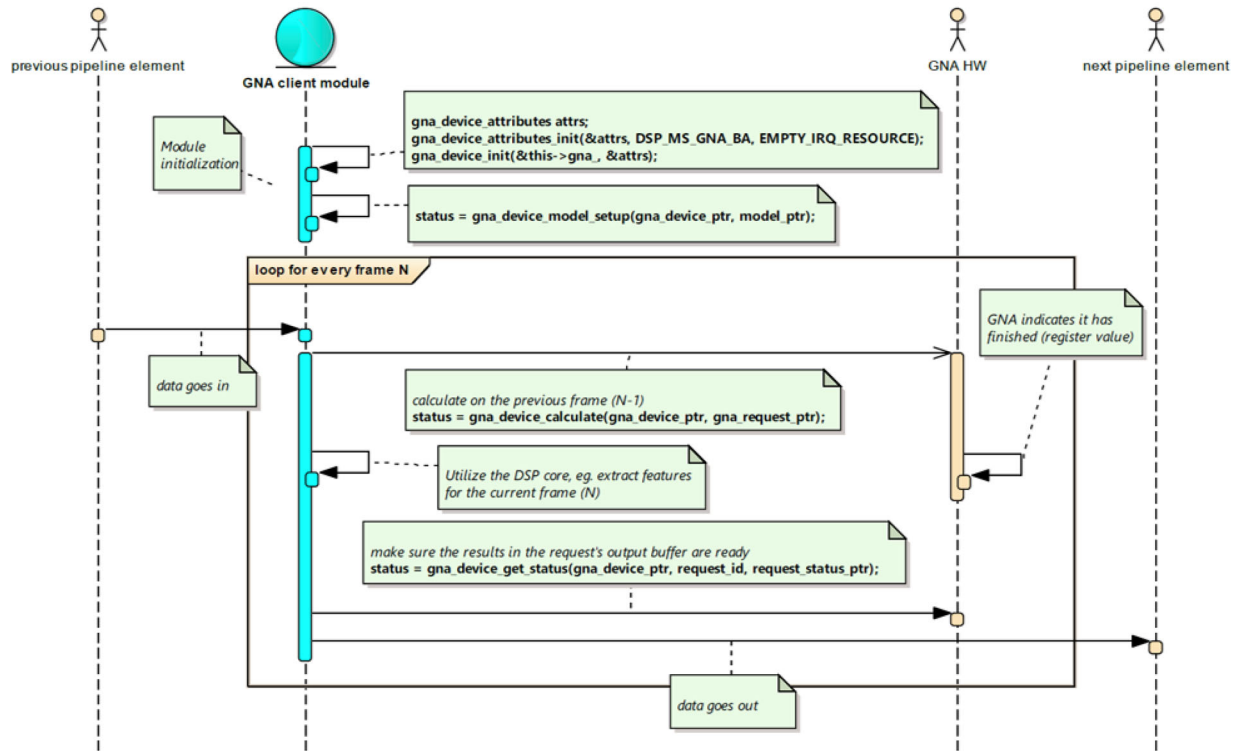

226.   The Intel GNA is a "flexible, low-power, high performance streaming co-processor that offloads neural network inference operations . . . ."[86] As explained by Intel in the programming guide excerpted below, the GNA is capable of utilizing extracted features from the DSP core that

---

[84] *See* Intel Newsroom, *Intel Launches First 10th Gen Intel Core Processors: Redefining the Next Era of Laptop Experiences* (Aug. 1, 2019), https://newsroom.intel.com/news/intel-launches-first-10th-gen-intel-core-processors-redefining-next-era-laptop-experiences/#gs.iwvy00.

[85] *See, e.g.,* Apple MacBook Pro 13-Inch "Core i5" 2.0 2020 4 TB 3 Specs, https://eve-rymac.com/systems/apple/macbook_pro/specs/macbook-pro-core-i5-2.0-quad-core-13-2020-4-thunderbolt-3-ports-scissor-specs.html (identifying Intel® Core™ i5-1038NG7 Processor); Intel Product Specification, Intel® Core™ i5-1038NG7 Processor, https://ark.in-tel.com/content/www/us/en/ark/products/196594/intel-core-i51038ng7-processor-6m-cache-up-to-3-80-ghz.html (identifying GNA in "Advanced Technologies").

[86] *See* G. Stemmer et al., *Speech Recognition and Understanding on Hardware-Accelerated DSP*, Interspeech (2017).

are calculated for frame (N) while calculating the acoustic model for the previous frame (N-1),

such that the GNA is processing one frame behind the DSP.[87]

In other words, the Intel GNA processes feature vectors (i.e. data) associated with audio signals received from the microphones (via the DSP) in the 2020 MacBook Pro 2.0GHz and is, in turn, an acoustic coprocessor.

---

[87] *See* Intel® Quark™ SoC S1000 Firmware Programming Guide, https://www.in-tel.com/content/www/us/en/developer/articles/guide/firmware-programming-guide-s1000.html.

227.   As to the first limitation of claim 10 of the '789 patent, the Accused Products contain a first interface for receiving at least one feature vector. Specifically, the block diagram at right, shows the Intel GNA with an interface for receiving feature vectors from an audio signal, shown as the end point of arrows from data arrays and network parameter arrays to input data buffers designed to "hold inputs…of each layer."[88]
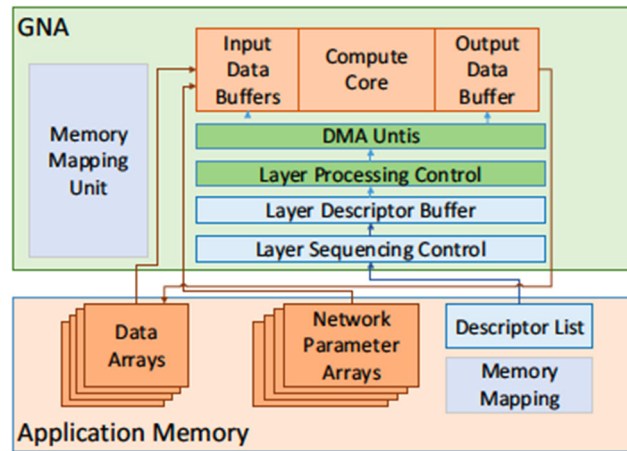

Figure 2 : *GNA block diagram and memory organization*
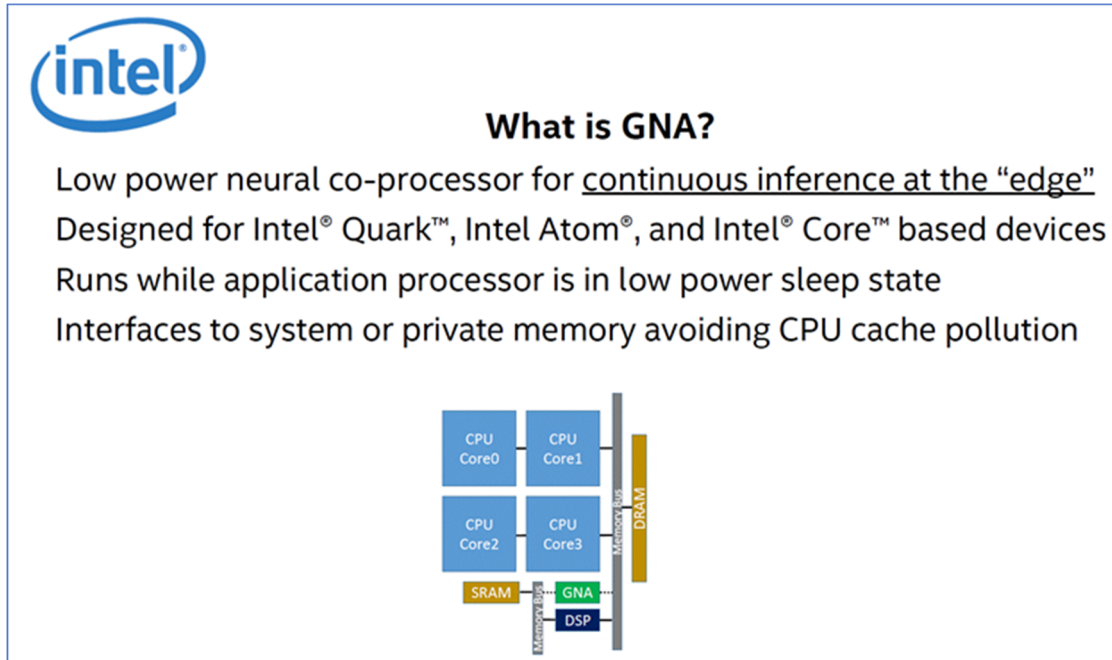
228.   The Accused Products also contain the second limitation of claim 10 of the '789 patent. On information and belief, the "Compute Core" of the Intel GNA, depicted in the figure above between the "Input Data Buffer" and "Output Data Buffer" is a calculating apparatus for calculating distances indicating a similarity between feature vectors and acoustic states of an acoustic model, read from an acoustic memory.

229.   At the 2017 International Conference of Acoustics, Speech, and Signal Processing (ICASSP), representatives of Intel demonstrated the GNA performing "continuous acoustic model likelihood scoring" using a 6-layer DNN with 2048 hidden nodes acoustic model. The demonstration showed the utility of the Intel GNA for highly accurate, low power, speech recognition compared to the application processor itself.[89]

---

[88] *See* G. Stemmer et al., *Speech Recognition and Understanding on Hardware-Accelerated DSP*, Interspeech (2017) at Figure 2.
[89] *See* Michael Deisher and Andrzej Polonsky, *Implementation of Efficient, Low Power Deep Neural Networks on Next-Generation Intel Client Platforms*, ICASSP (Apr. 12, 2017),

230.   Moreover, the compute core of the GNA retrieves the acoustic model from an

acoustic model memory, as shown below in a poster presented at ICASSP 2017 by Intel

employees.[90]

231.   The Accused Products further contain the third limitation of claim 10 of the '789

patent, a second interface for sending at least one distance calculated by the calculating apparatus.

Returning to the GNA block diagram and memory organization figure for the Intel GNA, a second

interface for sending distances calculated by the Compute Core is shown as the starting point of

an arrow from the output data buffer to the "Data Arrays."[91] The Output Data Buffers hold the

outputs of each layer.[92]

---

https://sigport.org/documents/implementation-efficient-low-power-deep-neural-networks-next-generation-intel-client.

[90] *See* Michael Deisher, et al., ICASSP GNA poster presentation "Implementation of Efficient, Low Power Deep Neural Networks on Next-Generation Intel Client Platforms," (ICASSP 2017), https://sigport.org/sites/all/modules/pubdlcnt/pubdlcnt.php?fid=1610 (noting that GNA can be programmed to compute a range of mathematical functions which are used to implement Neural Network based acoustic models, such as Affine, Convolutional, etc.).

[91] *Id.*

[92] *Id.*

232.　　　　Apple actively, knowingly, and intentionally induces infringement of one or more claims of the '789 patent under 35 U.S.C. § 271(b) by actively encouraging others to make, use, sell, and offer to sell in the United States, the Accused Products.

233.　　　　For example, Apple provides directions, instruction manuals, guides, and/or other materials that encourage and facilitate infringing use by others.

234.　　　　Apple has sold and is selling the Accused Products with the knowledge and intent that customers who buy the products will use the products for their infringing use and therefore that customers have been and are directly infringing the '789 patent.

235.　　　　Apple contributes to the infringement of one or more claims of the '789 patent under 35 U.S.C § 271(c) by offering to sell, selling, or importing into the United States, a component of the Accused Products that constitutes a material part of the inventions, knowing the same to be especially made or especially adapted for use in an infringement of the '789 patent, and is not a staple article or commodity of commerce suitable for substantial noninfringing use.

236.　　　　Apple was aware of the '789 patent at least from the date of this Complaint.

237.　　　　On information and belief, Apple has been aware of the '789 patent since its issuance on November 17, 2020.

238.　　　　On information and belief, Apple knew that the Accused Products infringe the '789 patent, or at a minimum believed there was a high probability that the Accused Products were covered by Zentian's patents, but willfully blinded itself to Zentian's patents and the infringing nature of the Accused Products.

239.　　　　The foregoing allegations are based on publicly available information and a reasonable investigation of the structure and operation of the Accused Products.

240.   Zentian reserves the right to modify this description of alleged infringement, including, for example, on the basis of information about the Accused Products that it obtains during discovery.

241.   The infringement details provided for claim 10 of the '789 patent are exemplary, and Zentian intends to assert additional claims of the '789 patent beyond claim 10.

242.   Apple's infringement has damaged and continues to damage Zentian in an amount to be determined at trial, but in no case less than a reasonable royalty for Apple's unauthorized use of the claimed inventions set forth in the '789 patent.

243.   This is an exceptional case. Zentian is entitled to attorneys' fees and costs under 35 U.S.C § 285 as a result of the infringement of the '789 patent by Apple.

DEMAND FOR A JURY TRIAL

244.   Zentian requests a trial by jury of all issues so triable under Federal Rule of Civil Procedure 38.

PRAYER FOR RELIEF

WHEREFORE, Zentian respectfully requests the following:

A.   That Judgment be entered that Apple has infringed one or more claims of each of the Asserted Patents, directly and indirectly, literally and/or under the doctrine of equivalents;

B.   An award of damages sufficient to compensate Zentian for Apple's infringement under 35 U.S.C. § 284, including an enhancement of damages on account of Apple's egregious willful infringement;

C.   That the case be found exceptional under 35 U.S.C. § 285 and that Zentian be awarded its reasonable attorneys' fees;

D.   Costs and expenses in this action;

E.   An award of prejudgment and post-judgment interest; and

F.   Such other and further relief as the Court may deem just and proper.

Dated: February 2, 2022   By: */s/ Jonathon K. Hance*
                                 Jonathon K. Hance (TX Bar No. 24065364)
                                 Lee Andrew Taggart (TX Bar No. 24093238)
                                 BRACEWELL LLP
                                 711 Louisiana St., Suite 2300
                                 Houston, Texas 77002
                                 Tel: (713) 223-2300
                                 Fax: (800) 404-3970
                                 jonathon.hance@bracewell.com
                                 drew.taggart@bracewell.com

                                 Nevin M. Gewertz (*pro hac vice* to be filed)
                                 Rebecca T. Horwitz (*pro hac vice* to be filed)
                                 BARTLIT BECK LLP
                                 54 West Hubbard Street
                                 Chicago, Illinois 60654
                                 Tel: (312) 494-4400
                                 Fax: (312) 494-4440
                                 nevin.gewertz@bartlitbeck.com
                                 rebecca.horwitz@bartlitbeck.com

                                 John M. Hughes (*pro hac vice* to be filed)
                                 Katherine L.I. Hacker (TX Bar No. 24074338)
                                 BARTLIT BECK LLP
                                 1801 Wewatta Street, Suite 1200
                                 Denver, Colorado 80202
                                 Tel: (303) 592-3100
                                 Fax: (303) 592-3140
                                 john.hughes@bartlitbeck.com
                                 kat.hacker@bartlitbeck.com

                                 Attorneys for Plaintiff Zentian Limited